

Essays on Worker Self Selection and Wage Inequality

Inaugural-Dissertation

zur Erlangung des Grades eines Doktors
der Wirtschafts- und Gesellschaftswissenschaften

durch

die Rechts- und Staatswissenschaftliche Fakultät der
Rheinischen Friedrich-Wilhelms-Universität Bonn

vorgelegt von

Felix Schran

aus Winterberg

Bonn, 2019

Dekan:	Prof. Dr. Jürgen von Hagen
Erstreferent:	Prof. Dr. Hans-Martin von Gaudecker
Zweitreferent:	JProf. Dr. Michael Böhm
Tag der mündlichen Prüfung:	27. September 2019

Für meine Eltern

Acknowledgements

Writing this thesis was both challenging as well as rewarding. It would not have been possible without the help of many people. First and foremost, I thank my supervisors and coauthors Hans-Martin von Gaudecker and Michael Böhm for their support and guidance in empirical work. It has been inspiring to learn from them, and it was a pleasure to work with them.

This thesis was carried out at several different institutes. I am grateful to the Bonn Graduate School of Economics (BGSE), the Institute of Applied Microeconomics (IAME), and the Institute of Labor Economics (IZA), in particular to the continuous efforts of Britta Altenburg, Simone Jost-Mussehl, Silke Kinzig, Benny Moldovanu, Andrea Reykers and Urs Schweizer for providing a great research environment. I am also grateful for financial support received from the German Science Foundation. Furthermore, I thank Christian Dustmann and the team at CReAM for making my research visit at UCL possible.

A big thank you goes to my friends and colleagues at the BGSE, in particular Zvonimir Bašić, Dmitry Kuvshinov, Jonas Radbruch, Louis Strang and Kaspar Zimmermann. Special thanks go to the BGSE football group and Holger Gerhardt for organizing our weekly matches.

Above all, however, I thank my family and especially my parents and my girlfriend Sarah for their unconditional faith and support. I dedicate my dissertation to them.

Contents

Acknowledgements	v
List of Figures	xiii
List of Tables	xvi
Introduction	i
References	iv
1 Occupation Growth, Skill Prices, and Wage Inequality	1
1.1 Introduction	1
1.2 Data, Literature, and Stylized Facts	3
1.2.1 Wage Inequality and Changes in Occupational Employment	4
1.2.2 Individual-Level Wage Growth and Selection	7
1.3 Estimating Skill Prices under Optimal Occupational Choice	10
1.3.1 A Tractable Model of Sector Choice	10
1.3.2 Price and Skill Changes	12
1.3.3 Estimation of the Model and Interpretation of Coefficients	15
1.3.4 Performance of the Estimation Strategy and Model Ex-	
tensions	18
1.4 Skill Prices and Skill Selection	20
1.4.1 Estimated Skill Price Changes and Skill Accumulation	21
1.4.2 Accounting for Skill Selection	25
1.4.3 Robustness of Results and the Task Content of Occupations	32
1.5 Skill Prices and Wage Inequality	35
1.5.1 The Attenuating Effect of Selection On Inequality	36
1.5.2 Factors Contributing to Wage Inequality	39
1.6 Discussion and Conclusion	43
1.A Further Details on the Data and Empirical Facts	45
1.A.1 Dataset Construction	45
1.A.1.1 Sociodemographics	45
1.A.1.2 Wages and Wage Growth	50
1.A.1.3 Sample Selection	51
1.A.1.4 Sample with Imputed Non-Employment Spells	51

1.A.2	Additional Stylized Facts	51
1.A.3	Constancy of Skill Accumulation	54
1.B	Theory	54
1.B.1	Proofs and Derivations	54
1.B.1.1	Derivation of Equation (1.3)	54
1.B.1.2	Derivation of Equation (1.5)	56
1.B.2	Multidimensional as Opposed to One-Dimensional Skills	56
1.B.2.1	Theory: Wage Gains When Switching Between Sectors	56
1.B.2.2	Evidence: Gains from Switching Into or Out of Ten Broad Occupations	58
1.B.3	Extensions of the Model	61
1.B.3.1	Costs of Switching Occupations	61
1.B.3.2	Non-Pecuniary Benefits	62
1.B.3.3	Learning About Skills	64
1.B.4	Occupation-Specific Fixed Effects as an Alternative Approach	65
1.B.4.1	Systematic Skill Accumulation	66
1.B.4.2	Idiosyncratic Skill Shocks	68
1.C	Monte Carlo Evidence	69
1.C.1	Data Generating Process	70
1.C.2	Key facts in the SIAB data	70
1.C.3	Baseline model	72
1.C.3.1	No Idiosyncratic Skill Shocks	73
1.C.3.2	Moderately Dispersed Shocks	76
1.C.3.3	Highly Dispersed Shocks	81
1.C.3.4	Persistent Shocks	86
1.C.4	Switching Costs	89
1.C.4.1	Benchmark: Moderate Switching Costs, No Shocks	90
1.C.4.2	Moderate Switching Costs, Moderate Shocks	92
1.C.4.3	High Switching Costs, Highly Dispersed Shocks	97
1.C.5	Changing Amenities	100
1.C.6	Occupation-Specific Fixed Effects	103
1.D	Additional Results for Section 1.4.1	105
1.D.1	Further Results of OLS Estimation	105
1.D.2	Instrumental Variable Estimates	107
1.E	Detailed Analysis of Skill Selection	109
1.E.1	Derivation of Equation (1.13)	109
1.E.2	Skill Changes and the Marginal Selection Effect	110
1.E.3	Sources of the Marginal Selection Effect	112
1.F	Robustness of Estimated Price and Skill Changes	114
1.F.1	Alternative Samples	114
1.F.2	Alternative Estimation Specifications	125

1.F.3	Task Measures and Changes in Occupations' Employment, Wages, Prices, and Skills	131
1.G	Further Details on Wage Inequality	131
1.G.1	Derivations and Further Results on the Attenuating Effect of Selection	131
1.G.2	Additional Results for the Scenarios from the Full Model	138
1.G.3	Effect of Skill Accumulation on Wage Percentiles	145
References	149
2	Changing Returns to Occupational Skill and Women's Wages	155
2.1	Introduction	155
2.2	Male and Female Wage and Employment Patterns	159
2.2.1	Data	159
2.2.2	Trends in Gender Wage Inequality and Employment Gaps	160
2.2.3	Life Cycle Employment and Wage Profiles	163
2.2.4	Distinguishing Changes in Average Wages from Changing Skill Prices	164
2.3	Impact of Changing Skill Returns on the Gender Wage Gap	167
2.3.1	Decomposition of the Average Gender Wage Gap	168
2.3.2	Proportion of Women Along the Wage Distribution	170
2.4	Conclusion	173
A	Additional Tables	175
B	Additional Figures	177
References	180
3	Locational Choice and Spatial Wage Inequality	183
3.1	Introduction	183
3.2	Connection to the Existing Literature	187
3.3	Spatial Changes in Wages and Employment	188
3.3.1	Data	188
3.3.2	The Geography of Wage and Employment Changes Across Germany	189
3.4	Estimating Local Wage Premia and Amenities	193
3.4.1	Model	193
3.4.2	Identifying Assumptions	196
3.5	Local Wage Premia and Worker Quality	199
3.5.1	The Role of a Shifting Occupation, Industry, and Education Structure	202
3.5.2	The Role of Local Amenities to Attract Skilled Workers	204
3.6	Sources of the Rising Density Wage Premium	205
3.7	Conclusion	208
A	Data Appendix	211
A.1	Aggregation of Local Labor Markets	211
A.2	Aggregation of Occupations	211

A.3	Aggregation of Industries	212
A.4	Education	212
A.5	Amenity Data	212
B	Observables of Growing Regions	213
B.1	Occupation Growth	213
B.2	Industry Growth	215
B.3	Educational Changes	217
B.4	Entrants' wages	219
C	Additional Results for Section 3.4.2	221
D	Additional Results for Section 3.5	223
E	Additional Results for Section 3.6	229
	References	231

List of Figures

1.1	Evolution of wage inequality and occupational employment . . .	6
1.2	Correlation of changes in employment, average wages, and wage growth	7
1.3	Selection into and out of occupations	8
1.4	Individual wage growth relative to 45–54 year olds	15
1.5	The evolution of skill prices	22
1.6	Skill accumulation of occupation stayers	23
1.7	Correlation of changes in employment, skill prices, and skills . .	24
1.8	Employment growth vs. the components of skill changes	27
1.9	Wage inequality scenarios	40
1.10	Selection into and out of occupations, controlling for age and education	52
1.11	Selection into and out of occupations, occupational switchers only	52
1.12	Changes in employment and average wages, 1975-2010	53
1.13	Wage growth between age groups	54
1.14	Gains from switching	59
1.15	Gains from switching (residual)	60
1.16	Descriptive statistics in the SIAB data	71
1.17	Descriptives, no shocks	74
1.18	Estimation results, no shocks	75
1.19	Descriptives, moderate shocks	77
1.20	Estimation results, moderate shocks	78
1.21	Descriptives, highly dispersed shocks	82
1.22	Estimation results, highly dispersed shocks	83
1.23	Descriptives, persistent shocks	87
1.24	Estimation results, persistent shocks	88
1.25	Descriptives, moderate switch costs, no shocks	90
1.26	Estimation results, moderate switch costs, no shocks	91
1.27	Descriptives, moderate switching costs and moderate shocks . .	93
1.28	Estimation results, moderate switching costs and moderate shocks	94
1.29	Descriptives, high switching costs and highly dispersed shocks .	98
1.30	Estimation results, high switching costs and highly dispersed shocks	99
1.31	Descriptives, trends in amenities	101

1.32	Estimation results, trends in amenities	102
1.33	Estimation results, no shocks as in Table 1.5	104
1.34	Estimation results, moderate shocks as in Table 1.6	105
1.35	Estimation results, highly dispersed shocks as in Table 1.9	105
1.36	The evolution of skill prices and skill accumulation of stayers	107
1.37	The evolution of average skill by occupation	110
1.38	Employment changes vs. accumulation and churning separately	111
1.39	Marginal selection of entrants and of leavers	111
1.40	Unemployment and leaving the labor force as a choice, i.e., filled non-employment spells	117
1.41	Including East Germans, foreigners, and women	121
1.42	Women only	122
1.43	Excluding anybody ever coded as a foreigner	123
1.44	All ages, 20–60 year olds	124
1.45	Flat spot identification using workers aged 45–54 years	127
1.46	Education-specific skill accumulation	128
1.47	Accounting for non-pecuniary benefits	129
1.48	Stint fixed effects estimation	130
1.49	Correlation of employment changes with task measures	132
1.50	Correlation of wage changes with task measures	133
1.51	Correlation of skill price changes with task measures	134
1.52	Correlation of skill changes with task measures	135
1.53	Wage inequality scenarios, anybody ever coded as foreign excluded	142
1.54	Wage inequality scenarios, filled non-employment spells	143
1.55	Wage inequality scenarios, order prices → accumulation → switch- ing	144
1.56	Skills accumulated during working life by percentile of wage dis- tribution	145
1.57	Shares in the wage distribution by quintile	147
2.1	Wage gap by cohort	164
2.2	Proportion of women in professions by cohort	165
2.3	Correlating employment, wage, and skill price changes	166
2.4	Wage gap by cohort including part-time workers	177
2.5	Proportion of women in professions by cohort	178
2.5	Proportion of women in professions by cohort	179
2.6	Profession × gender combinations along the wage distribution	179
3.1	Changes in employment and wages across West German regions	184
3.2	Relation between employment growth and wage growth	190
3.3	Wage growth differences	197
3.4	Worker quality increased in growing regions	200
3.5	Geography of changes in wage premia and skills, 1985 – 2010	201
3.6	Reweighted skill changes	204

3.7	Density wage premium increased over time	206
3.8	Changes in wage premia and worker quality by 1985's density	207
3.9	Changes in employment and wages by 1985's density	208
3.10	Spatial distribution of rural and urban Regions	211
3.11	Changes of occupational employment	213
3.12	Changes of occupational employment across space	214
3.13	Changes of industrial employment	215
3.14	Changes of industrial employment across space	216
3.15	Changes of education shares	217
3.16	Changes of education shares across space	218
3.17	Wages of regional entrants minus wages of incumbents	219
3.18	Entrants' relative wages across space	220
3.19	Wage growth between age groups	221
3.20	Changes in employment and wages including base period	222
3.21	Skill accumulation	223
3.22	Omitting controls for changes in amenities	224
3.23	Education dependent skill accumulation function	224
3.24	Occupation dependent skill accumulation function	225
3.25	Fixed effects estimation	226
3.26	Rewighted skill changes, education and industry	226
3.27	Amenities across space	227
3.28	Estimated amenities $\bar{\psi}_{l,2010}$ and the external amenity index	228
3.29	Relation between amenities and density	229
3.30	Wage premia vs density, omitting amenities, fixed effects	230

List of Tables

1.1	Contributions to marginal selection by origin and destination activities	29
1.2	Contributions to marginal selection by source of skills	31
1.3	Decomposition of the between-variance of wages, data and counterfactuals	37
1.4	Grouping of occupations	46
1.4	Grouping of occupations	47
1.4	Grouping of occupations	48
1.4	Grouping of occupations	49
1.5	Parameters	73
1.6	Parameters	76
1.7	True and estimated skill accumulation parameters, saturated OLS	79
1.8	True and estimated skill accumulation parameters, IV	80
1.9	Parameters	81
1.10	True and estimated skill accumulation parameters, saturated OLS	84
1.11	True and estimated skill accumulation parameters, IV	85
1.12	Parameters	86
1.13	Parameters	90
1.14	Parameters	92
1.15	True and estimated skill accumulation parameters, saturated OLS	95
1.16	True and estimated skill accumulation parameters, IV	96
1.17	Parameters	97
1.18	Parameters	100
1.19	Estimated skill accumulation coefficients (occupation groups, OLS)	106
1.20	Estimated skill accumulation coefficients (occupation groups, IV)	108
1.21	Estimated skill accumulation coefficients (occupation groups, OLS), filled non-employment spells	118
1.22	Contributions to marginal selection by origin and destination activities, filled non-employment spells	119
1.23	Contributions to marginal selection by source of skills, filled non-employment spells	120
1.24	Between-within occupation variance of observed log wages, experiments	137

1.25	Levels of wage percentiles and the variance for model and sample specifications	141
2.1	Gender gaps in wages and occupational employment	160
2.2	Wage and employment gaps when including part-time workers	162
2.3	Effect of changing skill prices on gender wage inequality	169
2.4	Decomposition of proportion of women by percentile	171
2.5	Gender wage gap decomposition by professions	175
2.6	Average gender wage gap decomposition including part-time workers	176
3.1	Observable determinants of changing worker quality	203

Introduction

From the end of the 1970s, inequality increased significantly in almost all industrialized countries. This trend has produced a huge economic literature examining possible causes of this rise. The most prominent explanations range from technological change to trade and globalization (Katz and Autor, 1999; Acemoglu and Autor, 2011).

The purpose of this thesis is to further contribute towards our understanding of increasing disparities in the labor market. For that, I investigate three different aspects of wage inequality: Chapter 1 analyses the influence of changing returns to skill on overall wage inequality – in the face of worker self selection. In Chapter 2, I study the importance of shifts in skill prices for declining wage differences between men and women. Last, Chapter 3 investigates the impact of worker sorting across locations on growing spatial disparities.

Apart from the topical focus on wage inequality, there are two further unifying aspects which connect all three chapters. First, the unit under study in all parts of this dissertation is the German labor market; partly because Germany allows the researcher to derive results from high quality, large scale administrative data which contain information on individual workers' behavior and associated outcomes: the "Sample of Integrated Labor Market Biographies".¹ However, Germany is also an interesting case to look at because of its impressive "resurgence" during the last two decades in terms of economic growth (Dustmann, Fitzenberger, et al., 2014) as well as the accompanied increase in wage inequality (Dustmann, Ludsteck, et al., 2009).

Second, all chapters are based on the idea that workers within an observed entity (e.g., occupation, firm, region, educational track) represent a self selected group as they make decisions based on their comparative advantage (Roy, 1951). On the one hand, workers' selection response raises serious challenges for identifying causes of increasing wage inequality within observational data. On the other hand, self selection represents an interesting object to study itself as workers' decisions may work against forces which essentially contribute to rising inequality (Heckman and Honoré, 1990). In fact, this is one of the

¹All chapters contained in this thesis make use of the factually anonymous Sample of Integrated Labor Market Biographies (version 7514). Data access was provided via a Scientific Use File supplied by the Research Data Centre (FDZ) of the German Federal Employment Agency (BA) at the Institute for Employment Research (IAB) under contract Number 101357.

central findings of the first chapter.

In Chapter 1, which is joint work with Michael Böhm and Hans-Martin von Gaudecker, we make two contributions to the existing literature. One is methodological, one is empirical. Firstly, we develop a new method to separate changes in skill prices (which are likely to stem from changes in labor demand because of technology or trade, for instance) from individual skill changes within a Roy model setting. This method represents the underlying framework for the following analyses as well as the approaches taken in Chapters 2 and 3.

Secondly, we apply our method to provide a more comprehensive picture of two of the most important trends in developed countries' labor markets during the last decades: a strong increase in wage inequality and a substantial reallocation of employment across occupations broadly characterized by polarization, i.e., a hollowing out of the middle class. More specifically, we study the relationship between changes in occupational employment, occupational wages, and rising overall wage inequality. Using long-running administrative panel data with detailed occupation codes, we first document that in all occupations, entrants and leavers earn lower wages than stayers. This empirical fact suggests substantial skill selection effects that are negative for growing occupations and positive for shrinking ones. We develop and estimate a model for prices paid per unit of skill in occupations, which incorporates occupation-specific skill accumulation over the career and endogenous switching across many occupations. Our results shed light on two important puzzles in prior literature. First, consistent with leading explanations for occupational employment changes, price and employment growth are positively related. Strong counteracting skill changes along the lines of our new empirical fact explain why occupational wages are unrelated to employment growth. Second, skill prices establish a long-suspected quantitative connection between occupational changes and the surge in wage inequality.

Chapter 2 utilizes the estimated skill price changes to shed light on the decline in wage inequality between men and women. Despite the findings from Chapter 1 suggest that changing occupational skill prices have strongly contributed to rising overall inequality, changing prices might have reduced wage inequality between men and women. Despite sounding paradoxical, the reason for that is simple: men and women work in very different occupations. This means, gender wage inequality might decline if skill prices rise within occupations important for female employment. Overall, I find that roughly 65% of the observed decline in the average gender wage gap between 1985 and 2010 can be explained by a reduction in occupational segregation between men and women. The remaining 35% are explained by shifts in occupational wages which increased within occupations important for female employment; and declined in many occupations important for male employment such as producing occupations. Motivated by the finding from Chapter 1 that average wages do not move as much as skill prices, though, I reestimate the part of the declining wage gap attributed to changes in (selection corrected) skill prices. The impact

of movements in these prices on the reduction in gender wage inequality has roughly been 13 percentage points larger than the impact of changes in average wages alone. Similar findings hold when decomposing the rise in the proportion of women at higher percentiles of the wage distribution and vice versa for lower percentiles. From a methodological point of view, this underscores the importance of accounting for selection effects in decompositions.

Last, Chapter 3 moves to the geographic aspect of inequality. During the last few decades, aggregate wage growth has been very unevenly distributed across space in Germany and other countries (Moretti, 2012). While wages in Southern German local labor markets rose by up to 28 log points, they increased only modestly or even declined in the north. Similar results apply to employment changes. Overall, this has led to a strong positive correlation between local wage and employment growth. What is driving these differential trends across space? Chapter 3 examines to what extent regions with growing employment are increasingly paying workers higher wage premia or, in contrast, to what extent the quality of workers in growing regions has risen. To decouple demand and supply for skill from each other, I estimate how regional wage premia paid for a unit of skill have changed over time using administrative panel data that allow me to hold constant changes in unobserved worker quality. I find that wage premia in regions with expanding employment did not rise more than in regions with declining employment. Instead, the quality of workers in growing regions went up. I investigate the importance of various possible observables for this relationship including local amenity differences, changes in occupation and industry structure as well as variation in education rates. Last, I explore the impact of changing wage premia and worker sorting on the recent rise of the density wage premium.

In summary, this dissertation consists of three self contained essays about the influence of worker self selection on the rise in wage inequality. Each chapter examines a different aspect of changing inequality with the purpose to get a better understanding of why we are increasingly living in a world of growing disparities. Comprehensive knowledge about the causes of inequality is essential, not least to guide policy for moderating future developments.

References

- Acemoglu, Daron and David H. Autor (2011). "Skills, Tasks and Technologies: Implications for Employment and Earnings". In: *Handbook of Labor Economics*, vol. 4 Part B. Edited by David Card and Orley Ashenfelter. Elsevier. Chapter 12, pages 1043–1171.
- Dustmann, Christian, Bernd Fitzenberger, Uta Schönberg, and Alexandra Spitz-Oener (2014). "From Sick Man of Europe to Economic Superstar: Germany's Resurgent Economy". In: *Journal of Economic Perspectives* 28.1, pages 167–188.
- Dustmann, Christian, Johannes Ludsteck, and Uta Schönberg (2009). "Revisiting the German wage structure". In: *Quarterly Journal of Economics* 124.2, pages 843–881.
- Heckman, James J. and Bo E. Honoré (1990). "The Empirical Content of the Roy Model". In: *Econometrica* 58, pages 1121–1149.
- Katz, Lawrence F. and David H. Autor (1999). "Changes in the Wage Structure and Earnings Inequality". In: *Handbook of Labor Economics*. Volume 3, pages 1463–1555.
- Moretti, Enrico (2012). *The New Geography of Jobs*. New York: Houghton Mifflin Harcourt.
- Roy, A. D. (1951). "Some Thoughts on the Distribution of Earnings". In: *Oxford Economic Papers* 3.2, pages 135–146.

Chapter 1

Occupation Growth, Skill Prices, and Wage Inequality

Joint with Michael Böhm and Hans-Martin von Gaudecker

1.1 Introduction

During the past decades, occupational employment has changed profoundly across Europe and the United States. A burgeoning literature has established fundamental shifts in labor demand as the most important cause of these changes (Autor, Levy, et al., 2003; Goos and Manning, 2007; Acemoglu and Autor, 2011; Goos, Manning, and Salomons, 2014).¹ Yet, it remains puzzling that neither occupational wages nor wage inequality show a clear reflection of these demand shifts. First, occupational employment growth has been decoupled from occupational wage growth (Goos and Manning, 2007; Mishel et al., 2013; Green and Sand, 2015; Hsieh et al., 2019; Taber and Roys, 2019). Second, while wage inequality has risen dramatically over the same period that occupational employment has changed, there remains debate in the literature about how much of this can be attributed to demand shifts (Autor, Katz, et al., 2008; Dustmann, Ludsteck, et al., 2009; Card et al., 2013; Firpo et al., 2013; Autor, 2019).

To solve these puzzles, we develop and estimate a model in which workers have occupation-specific skills that evolve endogenously over the career. Workers' optimal choices lead growing occupations to attract less skilled workers, which depresses these occupations' average wages. Shrinking occupations retain the most skilled parts of their workforce, lifting their average wages. The key distinction we make is between wages paid per constant unit of skill (*skill prices*)—which are directly affected by shifts in demand—and average occupa-

¹We discuss the literature in detail in the next section.

tional wages. Worker selection drives a wedge between the two. The evolution of skill prices will not be fully reflected in occupational wages and could even be neutralized or turn in the opposite direction. Moreover, between-occupation inequality will underestimate the impact of shifting occupational demand on wage inequality.

These mechanisms are consistent with stylized facts from our rich administrative panel data. As in the studies cited above, occupational wage growth and occupational employment growth bear no systematic relation with each other. At the same time, we find that individual workers' wage growth is substantially faster *within* expanding occupations. The discrepancy must stem from marginal workers. We newly document that workers who enter *any* occupation earn substantially less than incumbents. The same is true for workers leaving *any* occupation compared to stayers. Both effects are increasing in net occupation growth. The raw data thus reveal that net growth of an occupation will have a direct attenuating impact on average wages with selection operating in both directions. Growing occupations attract workers at the start of their careers, dragging down average wages. Declining occupations tend to shed workers who earn below-average wages, raising these occupations' average wages.

To quantify these effects, our economic model distinguishes between skill prices and skills. Workers have multidimensional skills that evolve heterogeneously across occupations and over the career. The model is explicitly based on Roy (1951) and therefore relaxes the exogenous mobility assumption (e.g., Abowd et al., 1999; Card et al., 2013; Cortes, 2016). That is, the choice of occupation may be driven by contemporaneous unobservable shocks. We employ a linear approximation to obtain an empirical formulation that is transparent and straightforward to estimate, even in settings with a large number of occupations. Our key identifying assumption is temporal stability of the skill accumulation function, which generalizes prior approaches. Accounting for this skill accumulation, the estimator then exploits workers' varying wage growth within and across occupations over time to identify changes in skill prices.

Our empirical analysis uncovers three main findings. First, there is a clear positive relationship between the development of skill prices and employment growth at the level of detailed occupations. This indicates that demand shifts were indeed the dominant drivers of both occupational employment and *skill-constant* wages over the past decades. Characterizing occupations by their task intensities, we find that the patterns are in line with routine-biased technical change (RBTC) as one of the important drivers of occupational demand.² More

²Note that this paper does not measure occupational demand or supply shocks directly. We instead infer from the co-movements of quantities and prices that these are consistent with demand shocks. Forces of occupational demand may include RBTC and related technical changes (e.g., Autor, Levy, et al., 2003), international trade and offshoring (Autor, Dorn, and Hanson, 2013; Goos, Manning, and Salomons, 2014), transformation of the industry structure (Bárány and Siegel, 2018), changes in consumption patterns (Autor and Dorn, 2013; Mazzolari and Ragusa, 2013), social skills content (Deming, 2017), among others.

generally, the patterns are consistent with polarization, since employment and skill prices of broad occupation groups with high as well as low wages increased compared to mid-wage occupation groups.

The positive correlation of occupational employment with skill prices and the lack of a correlation with average wages means that skills must deteriorate in growing compared to shrinking occupations. Our second main finding is that these skill changes from the estimation are consistent with those implied by our new empirical fact: lower-earning workers' net entry into growing and their net exit out of shrinking occupations fully account for the negative correlation between skill changes and employment growth. We term this the *marginal selection effect*. Viewed through the lens of our model, it stems from both entrants and leavers possessing lower skills than stayers in any occupation. We exploit the longitudinal dimension of the data to show that marginal selection conforms with economic notions of the underlying selection effects: The skill differences of entrants and leavers compared to stayers consist of differences in endowments, skill accumulation, and endogenous switching (staying) of those workers who experience negative (positive) shocks during their stint in an occupation.

Our third main finding is that occupational changes have driven much of the increase in wage inequality over the past decades. We decompose the trends in the wage distribution using our estimated model and find that changing skill prices in particular were a key driver of inequality. This impact is muted when studying between-occupation wage inequality in the raw data because average occupational wages do not systematically vary with skill prices, as implied by our second finding. Via reweighting, we also exploit the changing demographic structure to approximate some of the shifts of skill supply to occupations. These would have further raised inequality between occupations, had it not been for the strong selection effects.

This paper is structured as follows. Next, we describe the German SIAB data that we employ, relate to prior literature, and present the stylized facts motivating the course of our subsequent analysis. In the third section, we develop the model and estimation strategy. Section 4 presents the results on the evolution of skill prices, dissects the marginal selection effect, and reports on extensive robustness checks. In Section 5, we examine the impact of skill prices and skill selection on rising wage inequality. The last section discusses our findings' relationship to labor market institutions and sketches directions for further research.

1.2 Data, Literature, and Stylized Facts

We use the Sample of Integrated Labor Market Biographies (SIAB) provided by the IAB Institute at the German Federal Employment Agency. The SIAB is a 2% random sample of administrative social security records from 1975 to 2014. It is representative of 80% of the German workforce and includes employ-

ees covered by social security, marginal part-time workers, benefit recipients, individuals officially registered as job-seeking, and those participating in active labor market programs. The SIAB excludes the self-employed, civil servants, and individuals performing military service. Most notably, it contains individuals' full employment histories including detailed data on wages, industries, and occupations along with socio-demographics such as age, gender, or the level of education. The data is exact to the day as employers need to notify the employment agency upon changes to the employment relationship.

In order to work with a homogeneous sample throughout, we restrict the main sample to German men aged 25 to 54 years who are working full-time in West Germany. See Appendix 1.A.1 for the reasons behind these choices and for details on the wider dataset construction. We will relax all of these restrictions in robustness checks. We transform the spell structure into a yearly panel by using the longest spell in any given year, adjusting wages appropriately for spells that do not last the entire year. Due to a cap on social security contributions, 12% of wages are right-censored at this ceiling; we follow imputation procedures in Dustmann, Ludsteck, et al. (2009) and Card et al. (2013). We inflate all wages to 2010 prices using the German consumer price index.

A key strength of the SIAB data is that it provides high-quality longitudinal information on workers' occupations. Until 2010, the SIAB Scientific Use File contains a consistent set of 120 occupations; we cannot use subsequent years because the classification changes drastically thereafter. Most of our analyses will be based on the raw 120 occupations. To ease interpretation, we also aggregate them into broader groups following Acemoglu and Autor (2011) and others. These comprise managers, professionals, and technicians (Mgr-Prof-Tech); sales and office workers (Sales-Office); production workers, operators and craftsmen (Prod-Op-Crafts); and workers in services and care occupations (Srvc-Care). See Table 1.4 for the mapping of detailed occupations into these groups.

1.2.1 Wage Inequality and Changes in Occupational Employment

Two of the most important trends in developed countries' labor markets over the past decades have been a strong increase in wage inequality and a substantial reallocation of employment across occupations broadly characterized by polarization (for a summary see Acemoglu and Autor, 2011). As documented by, e.g., Spitz-Oener (2006), Dustmann, Ludsteck, et al. (2009), Card et al. (2013), and Goos, Manning, and Salomons (2014), Germany is no exception to either phenomenon. Figure 1.1 reproduces both trends in our dataset.

Figure 1.1a shows the trends of wage percentiles over the 1985–2010 period normalized to zero in 1985; thereby reproducing Figure 1 in Card et al. (2013) up to the normalization, sample, and the percentiles. Inequality increased strongly and steadily both in the upper half, measured by the difference between the 85th and the 50th percentile of log wages, and in the lower half (50 – 15 dif-

ference). These trends have arguably led to a broader debate about inequality and opportunity,³ as well as reignited policy efforts with regard to living wages and minimum wage regulations. For example, Germany introduced a statutory minimum wage in 2015; substantial raises to it are a constant source of public debate. Similarly, U.S. localities are in the process of or have already implemented a \$ 15 minimum wage (e.g., Jardim et al., 2017), more than twice the nationwide minimum wage.

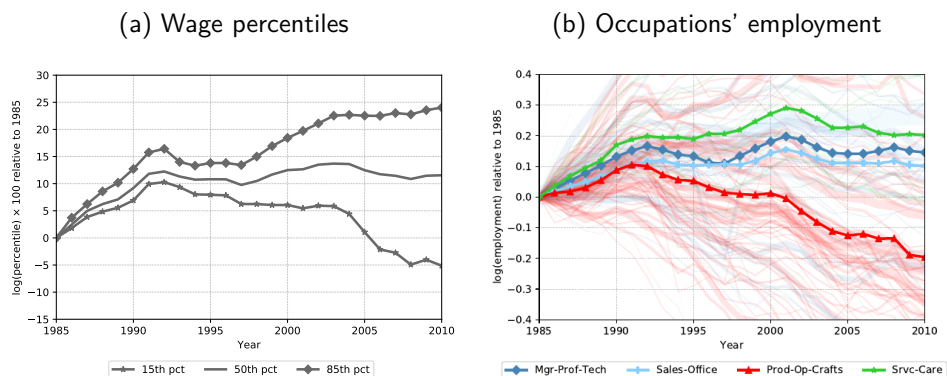
Using the year 1985 for the normalization once more, Figure 1.1b plots the trends in the logarithms of the detailed 120 occupations' employment (shaded lines) and the four aggregated groups (bold lines with markers). Employment in Production-Operators-Crafts occupations declined by more than 20 log points from a baseline share of over 60 percent, whereas the employment share of the other occupation groups increased. This trend has been termed "job" or "employment polarization" because Prod-Op-Crafts workers tend to be located in the middle of the occupational wage distribution (Goos and Manning, 2007). An important share of the declining employment in middle-paying occupations appears to be due to changes in technology (affecting codifiable routine-type jobs, see e.g., Autor, Levy, et al., 2003) as well as international trade and offshoring (affecting manufacturing-type jobs, e.g., Autor, Dorn, and Hanson, 2013). The resulting deterioration of employment opportunities—particularly severe for low and medium educated men—have been linked to societal trends of much wider concern.⁴

One may expect to see such shifts of the demand for different types of occupations directly in the wage distribution, not least because the wage and employment trends occurred largely in parallel (e.g., see Figure 1.1). There exists, however, surprisingly little quantitative evidence on the role of occupational change for the evolution of wage inequality: holding occupations' wages fixed at their initial levels and reweighting them with employment in subsequent decades, Goos and Manning (2007) show that composition effects due to employment polarization can account for a substantial part of changing wage inequality in the U.K. Very recently, Autor (2019) finds that in the U.S. a similar exercise explains only small shares of the income growth differentials across five education categories. Additionally accounting for the degree of urbanization comes close to matching the evolution of real wages of the non-college educated. For the German case, Dustmann, Ludsteck, et al. (2009) conclude that the rise of lower-half inequality was unlikely to stem from changes in demand. Card et al. (2013) run a set of Mincer regressions and incrementally add occupational identifiers, finding that the role of the latter for rising wage inequality is rather small.

³Among others, see the research agenda by Chetty et al. (e.g., 2011, 2018), which has also spilled over to Europe and Germany (e.g., Cornelissen et al., 2018).

⁴These include, among others, rising morbidity and mortality in midlife (Case and Deaton, 2015) as well as political polarization in various guises (Fetzer, forthcoming; Autor, Dorn, Hanson, and Majlesi, 2016).

Figure 1.1: Evolution of wage inequality and occupational employment

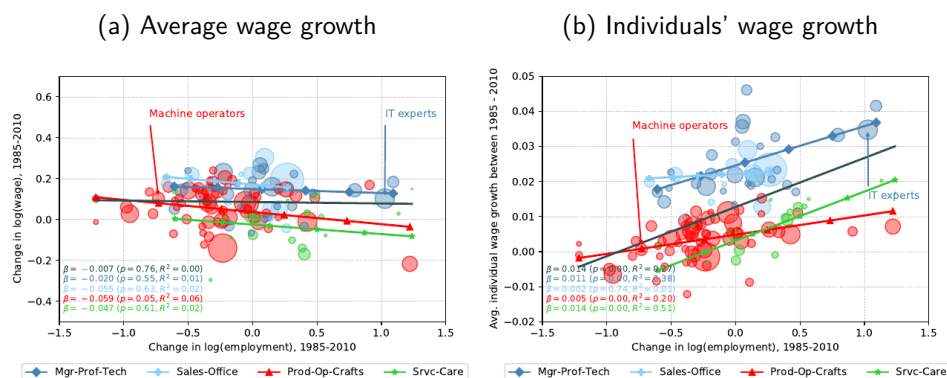


Notes: The vertical axis in Panel 1.1a shows the 15th, 50th and 85th log wage percentile over time relative to 1985. The vertical axis in Panel 1.1b shows the log change in the number of employed workers within an occupation over time. Shaded lines in the background represent the 120 detailed occupations in the SIAB SUF. The four groups show an aggregation of these detailed occupations as described in Appendix Table 1.4. The thickness of a shaded background line corresponds to the number of employed workers in an occupation averaged across years 1985 until 2010.

Figure 1.2a hints at why these types of analyses tend to have limited explanatory power. The graph plots changes in employment against changes of average wages over the 1985–2010 period for each of our 120 occupations. Variation along the horizontal axis shows that employment changes are very substantial. Many occupations grew or shrank by more than 50 log points. Yet, movements of average wages are surprisingly small given the variation in occupation growth and the large increase of wage inequality. Therefore, between-occupation decompositions—such as wage regressions with occupation dummies or reweighting strategies—may attribute little of the trends in wage inequality to factors like changing skill prices and employment structure, and much of its increase to unexplained within-occupation inequality. More fundamentally, the employment and wage changes in Figure 1.2a are uncorrelated: occupations that grew a lot did not experience larger average wage increases than shrinking occupations. To pick the two highlighted examples, IT experts' employment increased by 102 log points or 178% and their average wages grew by 10%, just above the overall average. Machine operators—a prototypical occupation one would expect to be negatively affected by routine-biased technical change—shrank by 73 log points or 51%. Yet, their average wages grew by the same amount as those of IT experts.

Within the broader groups, the non-correlation between wage and employment growth even turns negative for the lower-earning Prod-Op-Crafts and Srvc-Care occupations. This is consistent with the regressions reported by Dustmann, Ludsteck, et al. (2009) in their Section IV.D, which led them to conclude that demand shifts were unlikely to drive lower-end inequality. The finding of little or negative correlation between occupational wage and employment growth is not

Figure 1.2: Correlation of changes in employment, average wages, and wage growth



Notes: The vertical axis in Panel 1.2a shows the change in average wages between 1985 and 2010. The vertical axis in Panel 1.2b depicts individual wage growth averaged across years 1985 until 2010. The horizontal axis in both panels shows the change of the log number of employed workers within an occupation between 1985 and 2010. One bubble represents one of the 120 detailed occupations in the SIAB SUF. The four groups show an aggregation of these detailed occupations as described in Appendix Table 1.4. Bubble size corresponds to the number of employed workers in an occupation averaged across years 1985 until 2010. Regression lines across all occupations (black) and within the four broad groups (colored) are weighted by the number of employed workers.

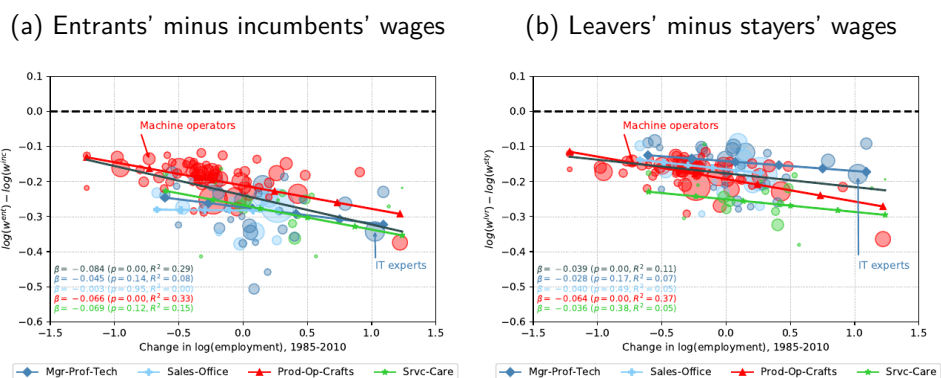
confined to Germany. Hsieh et al. (2019) and Taber and Roys (2019) document correlations between the growth rates of occupational employment and wages in the U.S. that are very small and positive or zero, respectively. Employment in low-skill occupations increased in the U.K. and Canada, while at the same time wages in these occupations dropped compared to routine occupations (Goos and Manning, 2007; Green and Sand, 2015). Next to the role that occupations have to play for wage inequality, this begets the more fundamental question of whether, on aggregate, shifts in demand versus supply of labor to different occupations were the dominant factor for the changes of the employment structure.⁵ We will find that, while the latter may have a role to play, the data strongly suggest that demand changes along the lines of routine-biased technological change or international trade are important.

1.2.2 Individual-Level Wage Growth and Selection

As a first pass, Figure 1.2b shows that there is a strong positive correlation between employment and *individual-level* wage growth. The horizontal axis is the same as in Panel a whereas the vertical axis plots the average annual wage growth of workers who stayed in their occupation for any two consecutive years. Wage growth rates within occupations clearly line up with their employment

⁵E.g., Glitz and Wissmann (2018) argue that a declining supply of medium versus low-skilled young workers in Germany was responsible for part of the rising lower-end inequality depicted in Figure 1.1a.

Figure 1.3: Selection into and out of occupations



Notes: The vertical axis in Panel 1.3a shows the average wage of an entrant to an occupation relative to the average wage of incumbents. The average is taken across years 1985 until 2010. The vertical axis in Panel 1.3b shows the average wage of a worker leaving an occupation next period relative to the average wage of stayers. The average is taken across years 1985 until 2009 to avoid all workers being leavers at the sample end. The horizontal axis in both panels shows the change of the log number of employed workers within an occupation between 1985 and 2010. One bubble represents one of the 120 detailed occupations in the SIAB SUF. The four groups show an aggregation of these detailed occupations as described in Appendix Table 1.4. Bubble size corresponds to the number of employed workers in an occupation averaged across years 1985 until 2010. Regression lines across all occupations (black) and within the four broad groups (colored) are weighted by the number of employed workers.

growth. Abstracting from other factors that we will control for later—most notably, occupation- and age-specific returns to experience—the main factor leading to the stark differences between the two panels of Figure 1.2 may well be differential selection into occupations. Put differently, demand shifts could indeed be driving the changes of employment and prices paid for skilled labor across occupations, but negative selection of entrants into growing occupations shrouds this relation when looking at average occupation-specific wages. The underlying occupational prices would be spreading out more than the average occupational wages, which are captured in the above-discussed decomposition analyses.

Data limitations have prevented a more thorough analysis of the presence and magnitude of such selection effects. In particular, the main sources in the U.S. are repeated cross-sections (CPS, Census) or longitudinal data too small in size for investigating individual-level dynamics across detailed occupations (PSID, NLSY). The SIAB data allow us to track occupational biographies over the entire career. Figure 1.3 gives more direct evidence on the importance of selection effects by plotting employment changes against the wage differentials between marginal workers who switch and inframarginal workers who stay in their occupations.⁶

⁶McLaughlin and Bills (2001) perform a related exercise with a coarser set of industry sectors in the PSID data. They report similar results on wage differences but struggle to find a correlation with changes in employment shares, possibly due to the small sample size.

The vertical axis of Figure 1.3a shows the difference between entrants and incumbents. An occupational entrant is defined as anybody who is newly observed in the occupation in the current period. He could be joining the labor force for the first time, switching from a different occupation, or entering from unemployment or outside of the labor force. The difference between this group and incumbents is strongly negative and strongly declining in occupation growth. The latter suggests that skill selection is the reason for the wage gaps—rather than, e.g., delayed wage contracting (Lazear, 1981)—since it is consistent with a situation where the skill pool that growing occupations can draw from shrinks with the extent of their expansion. Returning to our example, a machine operator might find it attractive to switch careers and become an IT expert in reaction to more lucrative employment opportunities there. It is not unreasonable, either, because controlling complex machines often involves some programming and algorithmic knowledge. However, given that he preferred a different career before, it is likely that the former machine operator’s specific skills are such that he will obtain a lower wage than incumbent IT experts.

In principle, the patterns in Figure 1.3a could be generated if occupation choice only happened at labor market entry in combination with substantial returns to experience. If this was the sole effect, however, we would expect that the wages of workers leaving their occupations would be higher than the wages of those who stay on. Put differently, in such a scenario individuals dropping out of our sample after age 54 should dominate the difference between leavers and stayers. Figure 1.3b shows that the opposite is the case. As for entrants, marginal workers have substantially lower wages than those who stay on. Again, the difference is increasing in employment growth. Put differently, only the lowest-skilled workers leave fast-growing occupations. All patterns in Figure 1.3 persist when controlling for age and education or considering only moves between occupations, i.e., discarding switches to or from non-working states.⁷ These pieces of evidence indicate that the wage gap is not just due to entrants being at an earlier stage of their career compared to incumbents.

The prominent models in the literature on occupational changes have difficulties matching Figure 1.3 because they feature one-dimensional skills (e.g., Acemoglu and Autor, 2011; Autor and Dorn, 2013). One-dimensionality ensures tractability in general equilibrium, which led to many important insights. The flipside is that it leads to a hierarchical ranking of occupations by skill, implying both that switchers to higher-ranked occupations leave lower-ranked occupations from above and that switchers from higher-ranked occupations enter lower-ranked occupations from above (Papageorgiou, 2014). This is hard to square with the fact that even entrants and leavers in low-wage occupations

⁷See the figures in Section 1.A.2 of the Appendix. When controlling for covariates, the magnitudes of the differences become smaller on average and the slopes tend to become more pronounced. As one would expect, considering only switches that happen directly between occupations has an attenuating effect on all patterns, but the qualitative pattern is always the same and highly significant.

generally earn less than incumbents and stayers, respectively.⁸

Instead, the patterns in Figure 1.3 call for a Roy-like approach to model sorting across occupations, with workers who possess specific skills such that both entrants and leavers are less skilled than incumbents and stayers. For example, in a model with two occupations and two skill types, Papageorgiou (2014) shows that switching workers earn wages below the average in the occupation they are leaving as well as the one they are entering so long as they do not have an absolute advantage in both occupations. Young (2014) calls this the case where “comparative advantage is aligned with absolute advantage”, which leads to declining skills in growing sectors. The conditions in Papageorgiou and Young are sufficient for marginal workers to have lower skills than inframarginal workers and for occupation growth causing skills to deteriorate. The necessary conditions are weaker; they only require that skills across occupations are not perfectly correlated and thus multidimensional (e.g., Heckman and Sedlacek, 1985). This level of generality forms our point of departure for the next section, where we develop a model that allows us to quantify these effects.

1.3 Estimating Skill Prices under Optimal Occupational Choice

This section presents our model to estimate skill prices, which enables us to distinguish price from selection effects when occupational wages change over time. We start by describing how we can exploit workers’ occupation choices in a classic Roy (1951) model to estimate the growth of potential wages across sectors. In Section 1.3.2, we outline our decomposition of wages into prices and workers’ skills along with a discussion of our main identifying assumptions. We then show how the model lends itself to a straightforward estimation strategy, which is feasible even for the 120 occupations \times 35 years in our application. In Section 1.3.4, we bring the estimation strategy to its limits in a series of Monte Carlo experiments and show how to incorporate additional features, for example, non-pecuniary job attributes in the generalized Roy model.

1.3.1 A Tractable Model of Sector Choice

There are $k = 1, \dots, K$ distinct occupations. At time t a worker i earns potential wages $W_{i,t} = (W_{1,i,t} \ W_{2,i,t} \ \dots \ W_{K,i,t})$. Most of our analysis will be in relative terms and we use lowercase letters to denote the logarithm of a variable. As in Roy (1951), we assume that workers maximize their incomes by

⁸Honing in on evidence similar to Figure 1.3, we explicitly test and reject the model of one-dimensional skills in Online Appendix 1.B.2. We do however obtain some evidence for a hierarchy between Mgr-Prof-Tech and the other broad occupation groups. This aspect of our data is consistent with the findings in Groes et al. (2014).

choosing the occupation in which they earn the highest wage:

$$w_{i,t} = \max\{w_{1,i,t}, \dots, w_{K,i,t}\} = \sum_{k=1}^K I_{k,i,t} w_{k,i,t}, \quad (1.1)$$

where $I_{k,i,t} \equiv \mathbb{1}[\max_{j=1,\dots,K}\{w_{j,i,t}\} = w_{k,i,t}] = \mathbb{1}[w_{k,i,t} \geq w_{j,i,t} \forall j \neq k]$ is a choice indicator for occupation k .

Ignoring the source of changes in potential wages in this subsection, we begin by considering the effect of marginal changes thereof on realized wages. By the envelope theorem,⁹ such changes will only have marginal effects on realized wages because occupation choices are the solution to the optimization problem (1.1). Put differently, workers do not enjoy discrete gains in realized wages when switching occupations in response to marginal changes of potential wages. For notational simplicity, we suppress the case of indifference at the prevailing wage (it will be trivially captured once we move to discrete wage changes immediately below) and write the marginal change in worker i 's realized wage at time t as:

$$dw_{i,t} = \begin{cases} dw_{1,i,t} & \text{if } I_{1,i,t} = 1 \\ \vdots & \\ dw_{K,i,t} & \text{if } I_{K,i,t} = 1 \end{cases}$$

or, equivalently,

$$dw_{i,t} = I_{1,i,t} dw_{1,i,t} + \dots + I_{K,i,t} dw_{K,i,t} = \sum_{k=1}^K I_{k,i,t} dw_{k,i,t}. \quad (1.2)$$

In order to understand wage changes between discrete time periods, we integrate over Equation (1.2) from potential wages $\{w_{1,i,t-1}, \dots, w_{K,i,t-1}\}$ to $\{w_{1,i,t}, \dots, w_{K,i,t}\}$. With a slight abuse of notation—made precise in Appendix 1.B.1.1—we obtain

$$\Delta w_{i,t} = \sum_{k=1}^K \int_{w_{k,i,t-1}}^{w_{k,i,t}} I_{k,i,\tau} dw_{k,i,\tau}. \quad (1.3)$$

This result is rather intuitive: if a worker stays in his occupation k' between two points in time ($I_{k',i,t-1} = I_{k',i,t} = 1$), his realized wage change is equal to the change in his potential wage in the chosen occupation (i.e., $\Delta w_{i,t} = \Delta w_{k',i,t}$). If the worker switches from some other occupation k'' to k' , ($I_{k'',i,t-1} = 1, I_{k',i,t} = 1$), his realized wage change is made up of two hypothetical components. One

⁹The optimization problem (1.1) fulfills the conditions for the general envelope Theorem 2 of Milgrom and Segal (2002). Incidentally, Milgrom and Segal derive the change in their general value function as an integral over the choices similar to our Equation (1.3). Böhm (2019) provides a derivation similar to ours in a static setting; he also provides special cases that do not even require the envelope theorem.

part stems from the wage change he would have experienced had he stayed in his previous occupation. The other part is the corresponding wage change had he been in the destination occupation all along. The relative size of both parts is determined by the point of indifference, i.e., the potential wages $w_{k',i,\tau^*} = w_{k'',i,\tau^*}$ so that $\Delta w_{i,t} = (w_{k',i,t} - w_{k',i,\tau^*}) + (w_{k'',i,\tau^*} - w_{k'',i,t-1})$. This trivially simplifies to $\Delta w_{i,t} = w_{k',i,t} - w_{k'',i,t-1}$, which is exactly the wage change that the definition of the realized wage (1.1) implies. The fact that only potential wages in his origin and destination occupations matter for the observed wage change makes sense given that the worker has comparative advantage in both of these occupations.

In empirical analyses, Equation (1.3) is directly observable for occupation stayers. That is, occupation choices on the right-hand-side and realized wage changes on the left-hand-side appear directly in the data. For switchers, we need to approximate the choices because we cannot observe switchers' point of indifference. We linearly interpolate the choice indicators for $\tau \in (t-1, t)$:

$$I_{k,i,\tau} \approx I_{k,i,t-1} + \frac{I_{k,i,t} - I_{k,i,t-1}}{w_{k,i,t} - w_{k,i,t-1}} (w_{k,i,\tau} - w_{k,i,t-1}) \quad (1.4)$$

Defining $\bar{I}_{k,i,t} \equiv \frac{1}{2}(I_{k,i,t} + I_{k,i,t-1})$ and combining Equations (1.3) and (1.4), we obtain

$$\Delta w_{i,t} = \sum_{k=1}^K \bar{I}_{k,i,t} \Delta w_{k,i,t}. \quad (1.5)$$

A detailed derivation is in Appendix 1.B.1.2. The intuition of Equation (1.5) after the approximation is the same as before: if a worker stays in his occupation, his wage gain is the change of his potential wage in that occupation. If the worker switches, he obtains equal parts of the origin and destination occupations' wage change. The strength of this result is that it allows to recover potential wage changes—even for switchers when occupational choice is endogenous—from panel data on occupation choices and realized wages, allowing for many occupations due to its simplicity. In particular, mean changes of potential wages can be recovered from a regression of first-differenced wages on “average” occupation choices $\{\bar{I}_{k,i,t}\}_{k=1}^K$. This hinges, of course, on the quality of the approximation. We first note that (1.4) is not an approximation at all for the majority of workers who stay in their occupation. To assess the impact of those who switch, we run a large set of Monte Carlo simulations. We will report on them in Section 1.3.4, noting here that the approximation in (1.4) is not a first-order concern.

1.3.2 Price and Skill Changes

We denote potential wages as the product of workers' skills

$S_{i,t} = (S_{1,i,t} \ S_{2,i,t} \ \dots \ S_{K,i,t})$ and the occupation-specific prices paid for

a unit of skilled labor $\Pi_t = (\Pi_{1,t} \ \Pi_{2,t} \ \dots \ \Pi_{K,t})$ that prevail in the economy.¹⁰ The worker's potential log wages become for all $k \in \{1, \dots, K\}$:

$$w_{k,i,t} = \pi_{k,t} + s_{k,i,t} \quad (1.6)$$

The framework outlined in the previous section relies on differences; we thus do not place any restrictions on the initial levels of prices or skills. The empirical challenge is to disentangle changes in prices from changes in skills. In order to do so, we impose some structure on the skill accumulation process, which we model by learning-by-doing on the job. Its speed is occupation-specific and depends on observables; working in one occupation k' impacts subsequent skills in all other occupations. In particular, we assume that for all $k \in \{1, \dots, K\}$:

$$\Delta s_{k,i,t} = \sum_{k'=1}^K I_{k',i,t-1} X'_{i,t-1} \Gamma_{k',k} + u_{k,i,t}. \quad (1.7)$$

The vector $X_{i,t-1}$ consists of a constant and observable variables controlling the speed of skill acquisition or depreciation via the vector $\Gamma_{k',k}$. Note that this formulation contains a full set of interactions of the skill accumulation coefficients $\Gamma_{k',k}$ with the covariates $X_{i,t-1}$. The summation term in (1.7) thus maps the previous occupation choice k' interacted with $X_{i,t-1}$ into skill changes in all potential occupations in the current period.

Our key identifying assumption is that the systematic part of the skill accumulation function (1.7) is time invariant. This is embodied in the fact that $\Gamma_{k',k}$ does not carry a time-subscript. Our condition is implied by an assumption made in virtually the entire literature studying occupational changes (e.g., Acemoglu and Autor, 2011; Firpo et al., 2013; Young, 2014; Gottschalk et al., 2015; Cortes, 2016; Bárány and Siegel, 2018; Yamaguchi, 2018; Böhm, 2019) that differences in returns to worker characteristics over time are due to changes in the returns to skills rather than changes in skill endowments. That assumption pins down the levels and growth rates of skills; ours does the same for the growth rates only. Our requirement is thus weaker in the sense that we do not place a restriction on the initial skill levels when entering an occupation¹¹ or on the precise contents of work within occupations, which may have changed (Spitz-Oener, 2006). We also richly model the skill accumulation function, most importantly including fully stratified occupation choices $I_{k',i,t-1}$ and

¹⁰Contrary to previous literature, we do not draw an explicit distinction between workers' skills and occupational tasks because we do not need the tasks as a dimension-reduction device. Our formulation is, however, perfectly general and nests, for example, the wage setting model in Firpo et al. (2013). We will eventually use tasks information to help interpret our results and connect back to previous work.

¹¹This is along with other papers using panel data (Cortes, 2016; Cavaglia and Etheridge, 2017). Removing the restriction on levels seems important in light of the first-order shifts in some observable characteristics. E.g., Carneiro and Lee (2011) show that in the U.S., the average skill of college graduates declined substantially as enrollment rates increased between 1960 and 2000. One may expect similar effects in Germany given that college completion rates doubled between the older and younger cohorts in our analysis.

ages $X_{i,t-1}$ among the observables, so that composition changes of workers' learning-by-doing are flexibly accounted for. Conditional on these observables, we do however assume that the speed of learning on the job has not changed over time. For example, a car mechanic in 2010 may well spend more time fixing electronics than his counterpart in 1975. A secretary will send e-mails rather than typing letters. But there is no temporal change in the speed at which these people get better at their jobs from one year to the next.

Our identifying assumption implies that within occupations, the ratio of wage growth across different groups remains constant over time. We check this in Figure 1.4, which plots the year-to-year wage growth of 25-34 year-olds (Panel 1.4a) and 35-44 year-olds (Panel 1.4b) minus the wage growth of 45-54 year-olds. We subtract the overall mean everywhere so that all eight lines should be flat at zero under our assumption. The lines in the right panel come very close to it. The left panel is somewhat noisier, particularly for the group of managers, professionals, and technicians. The noise is not too surprising given that many in this group enter the labor market at ages 25-34 and initial wage growth should be more susceptible to business cycles.¹² For example, the largest changes of wage growth across age groups can be found during the dotcom bubble in the late 1990s.¹³

Another identification strategy is to assume that seasoned workers' average skill growth is zero ("flat spot identification", Heckman, Lochner, et al., 1998), allowing us to interpret all occupation-specific wage growth over the decades as price growth (but still incorporating endogenous choices as derived in Equation (1.5)). We will explore this as a robustness check that yields qualitatively similar results to our main specification.

In terms of unobservables, we allow the joint distribution function $F(u_{1,i,t}, \dots, u_{K,i,t})$ to vary freely across occupations. For example, idiosyncratic skill shocks can be correlated among similar occupations in an unrestricted way. The restrictions we do impose are independence across individuals and an identical conditional distribution over time. That is, each skill shock's mean, conditional on all predetermined variables, is assumed to be zero:

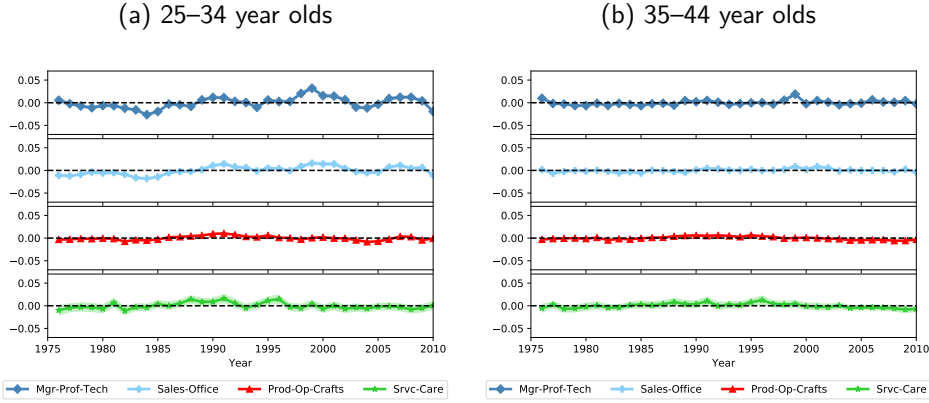
$$E[u_{k,i,t} | I_{k',i,t-1}, X_{i,t-1}] = 0 \quad \forall k', k \in K$$

These restrictions are considerably weaker than in existing fixed effects approaches (e.g., Abowd et al., 1999; Card et al., 2013; Cortes, 2016), which

¹²Consistent with that argument, Liu et al. (2016) find that the probability of an initial mismatch between jobs and workers is strongly countercyclical. This feeds into lower initial wages as well as persistently lower wage growth in subsequent periods.

¹³As an alternative for the 120 occupations, we split the sample in the middle (1993) and plot the change in log employment against the change in wage growth of young (age 25-34 or 35-44) minus old (45-54) workers in the resulting two periods. Naturally, there is more variation than for the four broad occupation groups but most of the occupations have very modest changes in relative wage growth rates and we cannot detect substantive patterns among them. See Online Appendix 1.A.3.

Figure 1.4: Individual wage growth relative to 45–54 year olds



Notes: The lines show average individual wage growth from $t-1$ to t by year of 25–34 (Figure 1.4a) and 35–44 (Figure 1.4b) year olds minus average wage growth of 45–54 year olds. Results are centered at zero to show trends over time. The shaded areas around the four lines are 95% confidence intervals. The four groups are based on an aggregation of detailed occupations in the SIAB SUF as described in Appendix Table 1.4.

require mean zero shocks conditional on *contemporaneous* variables ('exogenous mobility assumption'). We can do this because we explicitly derived wage growth when workers endogenously choose occupations in Section 1.3.1. We will see that this is important in the results below. Our restrictions are also more flexible than the types of assumptions that previous estimations of the Roy model (Heckman and Sedlacek, 1985) or of fully specified structural models (Lee and Wolpin, 2006) have invoked. In particular, we do not impose a parametric functional form for the distribution of unobservables.

1.3.3 Estimation of the Model and Interpretation of Coefficients

Under the assumptions we have made, we can compare price growth across different periods. The simplest intuition is that we can estimate the skill accumulation parameters $\Gamma_{k',k}$ in a base period $t = 0, \dots, T_{\text{base}}$ and use these to predict individuals' skill growth in $t = T_{\text{base}} + 1, \dots, T$ given their occupation choices. Subtracting predicted skill changes from realized wage growth and aggregating over all workers in an occupation yields price growth.

More formally, we substitute (1.7) into the equation for wage growth (1.5) to obtain our baseline estimation equation:

$$\Delta w_{i,t} = \sum_{k=1}^K \bar{I}_{k,i,t} \left(\Delta \pi_{k,t} + \sum_{k'=1}^K I_{k',i,t-1} X'_{i,t-1} \Gamma_{k',k} + u_{k,i,t} \right) \quad (1.8)$$

Our goal is to estimate the parameters in $\Delta \pi_{k,t}$ and $\Gamma_{k',k}$ for all $k, k' \in K$. As it stands, they are not separately identified from each other because of the intercept in $X_{i,t-1}$, which represents a level shifter for the speed of skill accumulation in each occupation by virtue of the interaction with last period's

occupational choice indicator. We can, however, compare the speed of skill price growth in different periods of our sample. Having to distinguish between price and skill growth is a general challenge of panel-data based estimations. We make the necessity for this explicit and will abstract from any short-term influences by using an entire decade (1975–1985) as the base period.¹⁴ Figure 1.12 in the Appendix depicts the employment and wage trends also for our base period. The decade 1975–1984 covers the entire business cycle. Between 1976 and 1979, average GDP growth was almost 4% annually; it then was below one percent on average until 1984 and picked up again in 1985. Furthermore, the resulting analysis period of 1985–2010 is the same as in Card et al. (2013).

In practice, we set $\Delta\pi_{k,t} = 0 \forall k \in \{1, \dots, K\}, t \in \{1, \dots, T_{\text{base}}\}$. The interpretation of $\Delta\pi_{k,t}$, $t \in \{T_{\text{base}} + 1, \dots, T\}$ changes depending on whether this holds as an assumption or whether it is better viewed as a normalization. The simplest interpretation obtains in the former case; i.e., skill prices during the base period were indeed constant. It is clear that the skill accumulation coefficients $\Gamma_{k',k}$ in this case will be identified from the base period. Accordingly, the estimates of $\Delta\pi$ can be interpreted as actual changes of skill prices for $t > T_{\text{base}}$.

Now suppose that constant skill prices during the base period are a poor approximation; i.e., there were substantial systematic changes between $t = 1$ and $t = T_{\text{base}}$. This implies that our estimated skill price changes in subsequent years are accelerations or decelerations relative to their (unknown) trends during the base period. To be precise, in the absence of other confounding factors, the estimated skill accumulation coefficients will be $\hat{\Gamma}_{k',k} = \Gamma_{k',k} + \frac{1}{2}\overline{\Delta\pi_{k,base}} + \frac{1}{2}\overline{\Delta\pi_{k',base}}$. Accordingly, the skill price estimates for $t = T_{\text{base}+1}, \dots, T$ identify $\Delta\hat{\pi}_{k,t} = \Delta\pi_{k,t} - \overline{\Delta\pi_{k,base}}$. In our discussion, we mainly stick with the easier literal interpretation of the parameter estimates. We will note the caveat on several occasions, taking particular care to point out instances where the acceleration/deceleration interpretation does make a substantive difference.

Turning to the estimation of the model, we first obtain a standard regression equation from (1.8) by writing out the summations:

$$\Delta w_{i,t} = \sum_{k=1}^K \bar{I}_{k,i,t} \Delta\pi_{k,t} + \sum_{k=1}^K \sum_{k'=1}^K \bar{I}_{k,i,t} I_{k',i,t-1} X'_{i,t-1} \Gamma_{k',k} + v_{i,t}, \quad (1.9)$$

where $v_{i,t} \equiv \sum_{k=1}^K \bar{I}_{k,i,t} u_{k,i,t}$. It is clear from this definition that v and the regressors are correlated since a large innovation to skills in a particular occupation makes it more likely that choosing this occupation happens to be optimal. First, we argue that a basic OLS regression of (1.9) will often yield good results. We then outline an instrumental variables strategy.

¹⁴As an alternative, Cortes (2016) and Cavaglia and Etheridge (2017) do not use a base period and thus implicitly set one of the skill accumulation parameters to zero (details in Appendix 1.B.4).

The regression (1.9) is a saturated skill model including all combinations of occupation choices $I_{k',i,t-1}$ and $I_{k,i,t}$. In the base period, the regression gives:

$$E \left[\Delta w_{i,t} \mid \{I_{k,i,t}, I_{k,i,t-1}\}_{k=1}^K, X_{i,t-1} \right] = E \left[\sum_{k=1}^K \sum_{k'=1}^K \bar{I}_{k,i,t} I_{k',i,t-1} X'_{i,t-1} \Gamma_{k',k} + v_{i,t} \mid \{I_{k,i,t}, I_{k,i,t-1}\}_{k=1}^K, X_{i,t-1} \right]$$

The fully interacted base period regression identifies this conditional expectation function and therefore yields expected skill changes:

$$E \left[\Delta s_{k,i,t} \mid \{I_{k,i,t}, I_{k,i,t-1}\}_{k=1}^K, X_{i,t-1} \right] \quad (1.10)$$

Defining $\nu_{i,t} \equiv \sum_{k=1}^K \bar{I}_{k,i,t} \left[\Delta s_{k,i,t} - E \left[\Delta s_{k,i,t} \mid \{I_{k,i,t}, I_{k,i,t-1}\}_{k=1}^K, X_{i,t-1} \right] \right]$, the regression equation in the analysis period can be re-written as:

$$\Delta w_{i,t} = \sum_{k=1}^K \bar{I}_{k,i,t} \Delta \pi_{k,t} + \sum_{k=1}^K \bar{I}_{k,i,t} E \left[\Delta s_{k,i,t} \mid \{I_{k,i,t}, I_{k,i,t-1}\}_{k=1}^K, X_{i,t-1} \right] + \nu_{i,t} \quad (1.11)$$

Conditional on $X_{i,t-1}$ and any combination of $I_{k',i,t-1}$ and $I_{k,i,t}$, the expectation of

$E \left[\Delta s_{k,i,t} - E \left[\Delta s_{k,i,t} \mid \{I_{k,i,t}, I_{k,i,t-1}\}_{k=1}^K, X_{i,t-1} \right] \right]$ is zero by construction. The

point here is that in the base period we already estimate the wage changes of occupation switchers, including the skill accumulation as well as idiosyncratic skill shocks. Therefore, if

$E \left[\Delta s_{k,i,t} \mid \{I_{k,i,t}, I_{k,i,t-1}\}_{k=1}^K, X_{i,t-1} \right]$ is consistently estimated in the base period, the error term in regression (1.11) is uncorrelated with the regressors $\bar{I}_{k,i,t}$ and changes in skill prices are identified under our assumptions.

An alternative approach to removing the bias in Equation (1.8) is by instrumenting the regressors $\{\bar{I}_{k,i,t}\}_{k=1}^K$ with their predetermined components $\{I_{k',i,t-1}\}_{k'=1}^K$, which are not a function of $u_{k,i,t}$. As in dynamic panel data models (Anderson and Hsiao, 1982; Arellano and Bond, 1991), we could in principle use long occupational histories as instruments. It is well-known, however, that this leads to issues with many weak instruments (e.g., Newey and Windmeijer, 2009). We thus instrument $\bar{I}_{k,i,t}$ to get $\pi_{k,t}$ with $I_{k,i,t-1}$, i.e., individual i 's occupation choice in the year before in order to have an instrument for skill price changes between years $t-1$ and t . For skill changes, we instrument $\bar{I}_{k,i,t} I_{k',i,t-1} X'_{i,t-1}$ with the occupational history in the two years preceding $t-1$, i.e., $I_{k,i,t-2} I_{k',i,t-1} X'_{i,t-1}$ and $I_{k,i,t-3} I_{k',i,t-1} X'_{i,t-1}$. This amounts to $(T - T_{base}) \cdot K + 2 \cdot K^2 \cdot L$ instruments, where L is the number of elements in X_t . This strategy will not be feasible for large K but we will use the IV as a major alternative specification for the four broad occupation groups.

Finally, notice that the OLS estimates $\hat{\Gamma}_{k',k}$ may not correspond to the structural skill accumulation parameters in Equation (1.9). The reason is that the $\hat{\Gamma}_{k',k}$ are the averages of skill changes, whether due to systematic accumulation or due to idiosyncratic shocks, of $k' \neq k$ switchers or $k' = k$ stayers. Since switching or staying is endogenous, we expect the skill accumulation parameters to be overestimated in the OLS. The IV does not have this problem and we expect skill accumulation estimates for stayers to be unbiased. But the first stage may be weak for predicting occupational switches and thus it may also be difficult in the IV to obtain the correct structural estimates of the off-diagonal elements in $\Gamma_{k',k}$.

1.3.4 Performance of the Estimation Strategy and Model Extensions

We test the limits of our estimation method in a broad range of Monte Carlo experiments, also exploring extensions of the underlying economic model. Furthermore, we compare the performance of our approach to an alternative that uses occupation-specific fixed effects pioneered by Cortes (2016). We limit ourselves to a short description of the results, all details can be found in Section 1.C of the Appendix.

In the Monte Carlo simulations, we aim to create a fairly realistic setting. We draw a sample of occupations and wages at labor market entry from our SIAB dataset. The remaining potential wages are drawn from truncated distributions so that the observed initial choice is optimal within the model. The subsequent trajectories of wages in all occupations are simulated using our estimates for price changes and skill accumulation, varying the dispersion of the idiosyncratic shocks across experiments. We stick to the four broad occupation groups and draw $100 \times 50,000$ careers for each experiment. This balances the ability to summarize the results on the one hand and broadly resembles the effective size of detailed occupations in our application on the other hand. Section 1.C.2 of the Appendix reports on some dimensions of the actual data—occupational switchers, the distribution of wage innovations, and the evolution of wage inequality—that serve as a backdrop for judging what may constitute reasonable values for simulation inputs like, for example, the variance of skill shocks.

In Section 1.C.3, we analyze the performance of our estimation method when the data generating process is precisely the one described in Sections 1.3.1–1.3.2. A detailed verbal description is provided at the beginning of 1.C.3; its four subsections contain tables and figures for varying specifications regarding the distribution of the idiosyncratic skill shocks. In order to judge the quality of the approximation (1.4), we first shut these shocks off altogether. The only randomness in this experiment comes from the initial draws and from the evolving prices at the aggregate level. None of the 4×100 estimated lines is visually discernible from the respective truth; we thus conclude that the

approximation of individual wage growth under optimal occupation choice in Equation (1.4) is unlikely to be causing a bias in our basic setting.

We then set the standard deviation of idiosyncratic skill shocks to half of the standard deviation of innovations to wages in the SIAB. This yields switching behavior, wage innovations, and an evolution of the wage structure very similar to those in that actual data; we thus term this distribution to have “moderate shocks”. As predicted at the end of the previous section, the OLS estimates show a modest upward bias of stayers’ skill accumulation coefficients, whereas the IV estimates are almost exactly on target.¹⁵ Both sets of skill price estimates track the evolution of their actual values very closely. Intuitively, mistakes we make with respect to the structural accumulation in the base period cancel out in the estimation period, i.e., in Equation (1.11) for the OLS. This basic pattern holds true even when tripling the size of the shocks.¹⁶ We overestimate skill accumulation, particularly when using OLS, but skill price estimates remain remarkably close to their targets. Finally, adding persistence to the skill shocks by means of an AR(1)-process does not alter these conclusions either.

One aspect that previous literature has emphasized are fixed costs of switching occupations (e.g., Cortes and Gallipoli, 2017). In our framework, the point of indifference between staying in an occupation and switching will now be determined by wages adjusted for switching costs. This means, however, that unadjusted wages of switchers will exhibit jumps at the indifference point, introducing a potential bias to our estimates. We work this case out theoretically in Section 1.B.3 of the Appendix; Section 1.C.4 presents Monte Carlo analyses examining the bias’ importance. First, in a model without skill shocks and with moderate switching costs (5% of annual wages), our approximation (1.4) continues to work well. OLS estimates recover skill prices and stayers’ skill accumulation coefficients almost exactly in such a specification. As previously, we then add moderate and large skill shocks, paired with moderate and high (20% of annual wages) switching costs. All pictures show that the basic conclusions from the corresponding exercises without switching costs remain the same: We slightly overestimate the structural skill accumulation coefficients,¹⁷ but skill prices are estimated with remarkable precision.

Another key extension of our approach is to the generalized Roy model, including non-pecuniary values of occupations in the worker’s decision problem (e.g., as in Lee and Wolpin, 2006). Similar to the case with switching costs, workers who move to an occupation with lower (higher) non-pecuniary value will exhibit positive (negative) jumps in wages to compensate for the amenity differ-

¹⁵Also as expected, the cross-accumulation parameters are generally upward-biased in the OLS; and in the IV with weak instruments, they are large in absolute values.

¹⁶The descriptives on the resulting data in 1.C.3.3 show that tripling the shocks is clearly an extreme case. There is far more switching in all directions compared to the SIAB, wage growth is twice as high and more dispersed than in the data, and wage inequality is skyrocketing.

¹⁷As one would expect based on our theoretical analysis, the inertia generated by switching costs leads to a somewhat larger overestimation of the off-diagonal elements of Γ .

ence. We show formally in Appendix 1.B.3 that, if the non-pecuniary values are time-constant, the skill accumulation parameter $\hat{\Gamma}_{k',k}$ in our main specification will absorb them. If they are time-changing, the estimation Equation (1.9) has to be augmented and include regressors for occupation switches ($\Delta I_{k,i,t}$) on top of average occupation choices ($\bar{I}_{k,i,t}$) to control for (and estimate) the respective “wage compensation”. Section 1.C.5 of the Monte Carlo analyses examines such a case with rising amenities in one of the occupations, finding that the $\Delta I_{k,i,t}$ correction is indeed necessary but then we recover the skill prices and skill accumulation as well as before (plus the changing amenities themselves).

We also show formally in the Appendix that what we have referred to as idiosyncratic skill shocks is observationally equivalent in our analysis to a basic model of employer learning about workers’ skills (e.g., as in Altonji and Pierret, 2001; Gibbons, Katz, et al., 2005). This is due to the fact that log-linearity allows us to write the model in terms of expected skills, which can evolve because of changes in actual skills (our formulation above) or because employers change their expectations about individuals’ skills over time. The two interpretations are not mutually exclusive, of course.

Finally, we examine an alternative panel data approach for estimating skill prices due to Cortes (2016), who uses individual \times occupation specific fixed effects in order to control for skill selection. First, we show theoretically how to generalize Cortes’ estimation in order to flexibly control for a rich model of worker skill accumulation. We then implement this approach in the Monte Carlo simulations and find that it performs well in most cases. Exceptions are specifications with a lot of switching (i.e., a large number of occupations K or large skill shocks), when the ‘exogenous mobility’ assumption of fixed effects approaches discussed in Section 1.3.2 and Appendix 1.B.4 becomes quantitatively important. We conclude that the generalized version of Cortes’ method is a useful alternative when the goal is to estimate low-dimensional skill prices; it seems less suitable for applications that feature a large number of occupations.

1.4 Skill Prices and Skill Selection

This section first presents the estimation results for our main model. These include the evolution of skill prices, the accumulation of skills over the career, and the relation of prices and occupations’ average skills with employment growth. We then dig deeper into the nature of the implied selection effects, showing that the skill differences between marginal workers and those who remain in their occupations drive the strongly negative association of employment growth with average skill changes in an occupation.

Throughout the section, we focus on the OLS results because they allow us to estimate the model for both the 120 detailed occupations and the four broad groups. We describe the IV results for the latter along the way. In our main specification, $X_{i,t-1}$ contains two dummies for age groups 25–34 and 35–44 in

$t - 1$ and an intercept representing the omitted age group 45–54 (recall from Section 1.3.2 that these are fully interacted with occupation choices). In the final part of this section, we show that our results are robust to a variety of alternative choices regarding data preparation, sample selection, and estimation specification before connecting our results to the literature using a task-based approach.


1.4.1 Estimated Skill Price Changes and Skill Accumulation

Figure 1.5 depicts the evolution of skill prices, normalizing them to zero in 1985 and cumulating the yearly changes until 2010. In the broad occupation groups, skill prices increased strongly among Mgr-Prof-Tech occupations, modestly among Sales-Office and Srvc-Care; they decreased among Prod-Op-Crafts. The thin lines in the background show that these broad estimates mask substantial heterogeneity among the 120 detailed occupations. We will explore this in greater detail below.

Several distinct periods are noticeable. All prices increased during the favorable economic conditions between 1985 and 1991, although this was already less pronounced for the Prod-Op-Crafts occupations. These have experienced a continuous decline thereafter to the point that prices in 2010 were more than five percent below their initial value in 1985. For the other occupations, there was a drop during the 1992–93 recession as well; prices then stayed constant until they rebounded before the turn of the century. This rebound was most pronounced for Mgr-Prof-Tech occupations; prices in this group did not change much for the remainder of our sample period. Skill prices fell by about 5 percentage points for Sales-Office and Srvc-Care occupations between 2000 and 2010. All these broad patterns also hold up in the instrumental variable estimates with slightly different numerical values; see Figure 1.36 in the Appendix. They are consistent with the job polarization of Figure 1.1b above; even the temporal changes of employment and skill prices seem to be broadly aligned in the four broad occupations. We will analyze in detail this relationship between employment and price changes, and in Section 1.4.3 we use tasks measures to approximate the role of RBTC versus other underlying factors for these patterns in the German labor market.

Figure 1.6 graphs the estimates of the skill accumulation parameters for stayers (i.e., $\Gamma_{k,k}$) in the four broad groups and for the 120 detailed occupations. Skill growth in the early years of the career is steep. Absent changes in skill prices, it implies a 20% growth in wages between age 25 and age 34 for Prod-Op-Crafts or Srvc-Care occupations and 50% or more for the other two. It slows down mid-career and flattens out or turns negative toward the end of prime age. This reflects the well-established concavity of life-cycle wage profiles (e.g., Lagakos et al., 2018). Skill growth differs substantially by occupation. It is initially very fast in high-earning Mgr-Prof-Tech and Sales-Office occupation groups and never completely ceases. Growth is flatter and eventually peters out

Figure 1.5: The evolution of skill prices



Notes: The figure shows estimated skill price changes. OLS estimates as described by Equation (1.9). Shaded lines in the background represent the 120 detailed occupations in the SIAB SUF. The four groups show an aggregation of these detailed occupations as described in Appendix Table 1.4. The thickness of a shaded background line corresponds to the number of employed workers in an occupation averaged across years 1985 until 2010. The shaded areas around the four lines are 95% confidence intervals.

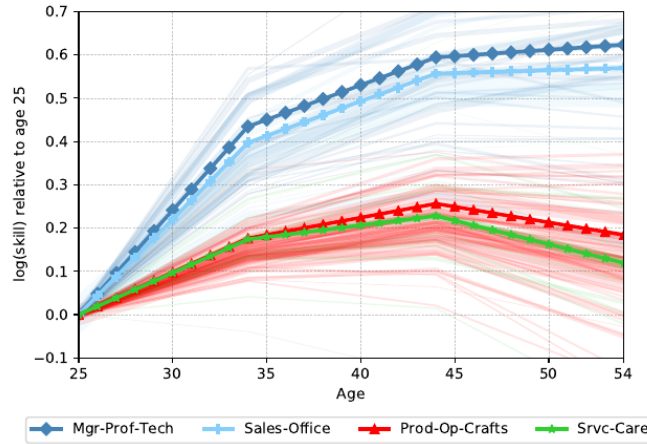
or turns negative in the Prod-Op-Crafts and Srvc-Care groups, i.e., occupations that often require physical labor.¹⁸ Again, the broad groups mask substantial heterogeneity across the detailed occupations. At the same time, it is the case that the “blue” occupations on the one hand and the “red and green” occupations on the other hand are almost separate; there are hardly any occupations to be found in the other block. This shows that life-cycle wage profiles are decidedly different across occupations and controlling for this fact is critical in producing reasonable estimates of prices and skills.

Our discussion in Sections 1.3.3–1.3.4 has shown that the parameter estimates depicted in Figure 1.6 may incorporate both the structural coefficients $\Gamma_{k,k}$ and shocks. The reason is dynamic selection: Stayers are more likely to have experienced favorable draws in their occupation,¹⁹ whereas IV estimates should show less such bias. As expected, the IV estimates for the four occupation groups are slightly lower, but none of the broad patterns change. The full set of our $\Gamma_{k',k}$ estimates for the four occupation groups can be found in Sec-

¹⁸One fact to note about the occupations in Germany compared to other countries is that Sales-Office is quite high-earning. Its average wages for men are about halfway between Mgr-Prof-Tech and Prod-Op-Crafts, employment is not declining over time, and we estimate rapid skill accumulation as well as rising skill prices for this occupation group. Using survey data, Cavaglia and Etheridge (2017) also document substantially higher wages for sales and office occupations in Germany than in the U.K.

¹⁹As noted before, for the estimation of skill prices this is a core strength of our approach because we allow for endogenous staying as well as switching of occupations.

Figure 1.6: Skill accumulation of occupation stayers



Notes: The figure shows estimates for stayers' skill accumulation during the life cycle. OLS estimates as described by Equation (1.9). Shaded lines in the background represent the 120 detailed occupations in the SIAB SUF. The four groups show an aggregation of these detailed occupations as described in Appendix Table 1.4. The thickness of a shaded background line corresponds to the number of employed workers in an occupation averaged across years 1985 until 2010. The shaded areas around the four lines are 95% confidence intervals.


tion 1.D of the Appendix. The OLS coefficient estimates show that switching into Mgr-Prof-Tech and Sales-Office goes in hand with substantial gains. The magnitude of the off-diagonal elements of Γ suggests that these incorporate sizeable idiosyncratic shocks.²⁰ The IV estimates for switchers also seem large in absolute value, which is not surprising given that the instruments can only weakly predict who switches out of his occupation. For the purposes of this paper, which requires controlling for, but not predicting, occupation switches, it is perfectly fine to identify the average gains associated with changing occupations. The critical task at hand for the skill accumulation function is to appropriately account for any kind of wage growth that may be due to observables or unobservables changing over the career; the OLS estimates do serve this purpose.

We now hone in on our key finding of this section, namely that employment growth and skill price growth go hand in hand. Figures 1.1b and 1.5 show that—consistent with shifts of occupational demand—both employment and skill prices in the broad Mgr-Prof-Tech, Sales-Office, and Srvc-Care groups increased compared to Prod-Op-Crafts. Figure 1.7a shows that detailed occupations' log employment changes between 1985 and 2010 are positively related to cumulated skill price changes over the same period. The upward-sloping

²⁰Judged against the difference between the true values and coefficient estimates in the Monte Carlo experiments of Section 1.3.4, they are in line with the specifications with “moderate” shocks.

black regression line summarizes this strong relationship for the 120 occupations, which is in marked contrast to the zero correlation for wages not corrected for composition effects (Figure 1.2a). As shown by the respective sub-regression lines, the relationship also holds within occupation groups. This indicates that our result is more general than a particular demand shifter that predominantly impacts broad occupation groups.

Figure 1.7: Correlation of changes in employment, skill prices, and skills



Notes: The vertical axis in Panel 1.7a shows the change in skill prices between 1985 and 2010 estimated with OLS as detailed in Section 1.3.3. The vertical axis in Panel 1.7b depicts the change in skills between 1985 and 2010 estimated as the residual between price and wage changes as shown in Equation (1.12). The horizontal axis in both panels shows the change of the log number of employed workers within an occupation between 1985 and 2010. One bubble represents one of the 120 detailed occupations in the SIAB SUF. The four groups show an aggregation of these detailed occupations as described in Appendix Table 1.4. Bubble size corresponds to the number of employed workers in an occupation averaged across years 1985 until 2010. Regression lines across all occupations (black) and within the four broad groups (colored) are weighted by the number of employed workers.

The deviations of the bubbles from the overall regression line can be informative about differences in labor supply to occupations. For example, Mgr-Prof-Tech occupations tend to be located to the right of the graph and above the grand regression line, meaning that prices have grown strongly compared to employment. This pattern is consistent with a combination of positive demand shifts and comparatively inelastic labor supply to those occupations, which seems plausible because of occupational licensing rules and often high educational requirements in Mgr-Prof-Tech. We estimate the largest increase in skill prices for physicians and pharmacists (bubble at the very top of Figure 1.7a), where educational requirements are high, places in medical school limited, and licensing rules very strict. In contrast, this argument does not apply to IT experts and the corresponding bubble is located below the overall regression line.

It is also possible to find examples where contemporaneous shifts of labor supply and demand seem to be important. For example, the right-most red bubble in Figure 1.7a are “assistants without further specification”. This occupation has arguably experienced a strong positive shock to labor supply, with many low-skilled immigrants and ethnic Germans from Eastern Europe entering

it after 1990. At the same time, temporary work agencies substantially increased their demand for this occupation. Taken together, this may generate the pattern that in Figure 1.7a, these assistants' skill prices remain almost constant while their average skills decline strongly according to Figure 1.7b. We will not attempt to distinguish different labor supply elasticities from contemporaneous shifts of supply and other factors. These are economic forces that may generate the variation around the regression line in Figure 1.7a; its positive slope indicates that demand shifts are the dominant force driving occupational changes.

Figure 1.7b depicts occupational employment growth against the cumulative changes of average skills implied by the skill price estimates. For every occupation, this is the difference between growth of its average wage (Figure 1.2a) and its skill price (Figure 1.7a), i.e., the second term on the right hand side of:

$$\underbrace{E[w_{i,t}|I_{k,i,t} = 1] - E[w_{i,t-1}|I_{k,i,t-1} = 1]}_{\text{Mean wage change}} = \quad (1.12)$$

$$\underbrace{\Delta\pi_{k,t}}_{\text{Price change}} + \underbrace{E[s_{k,i,t}|I_{k,i,t} = 1] - E[s_{k,i,t-1}|I_{k,i,t-1} = 1]}_{\text{Mean skill change}}$$

summed over the years $t = 1985, \dots, 2010$. The x-axis once more has occupations' growth over the analysis period. Figure 1.7b thus shows that implied skill changes constitute the flipside of the skill price estimates in the sense that growing occupations' decline of skills is strong. For example, the overall regression line indicates that average skills of the occupations that experience the fastest growth declined by 35 log points on average compared to those that shrank the most. These are large effects; we thus devote the next section to examining their components and their plausibility.

1.4.2 Accounting for Skill Selection

We have documented in Section 1.2.2 that entering (leaving) workers' skills on average appear decidedly below those of incumbents (stayers) and that faster-growing occupations draw even less skilled entrants (leavers). Given that growing sectors by definition experience net entry, this could substantially drag down growing occupations' average wages despite rising demand and increasing skill prices. Here we formalize and quantify this effect in the context of our model, showing that it indeed drives the systematic part of the relationship between employment growth and skills.

The change in average skills of an occupation in Equation (1.12) is determined by three mutually exclusive groups of workers: Those who leave the occupation after period $t - 1$; those who stay on after period $t - 1$ and are thus incumbent in period t ; and those who enter in period t . Denoting the share of leavers in $t - 1$ by $p_{k,t-1}^{lvr}$ and the share of period- t entrants by $p_{k,t}^{ent}$, we can

decompose the change of average skills in occupation k into three terms:

$$\begin{aligned}
 & \underbrace{\mathbb{E}[s_{k,i,t} | I_{k,i,t} = 1] - \mathbb{E}[s_{k,i,t-1} | I_{k,i,t-1} = 1]}_{\text{Mean skill change}} = \tag{1.13} \\
 & \underbrace{\left(1 - \frac{p_{k,t-1}^{lvr} + p_{k,t}^{ent}}{2}\right) \cdot \mathbb{E}[\Delta s_{k,i,t}^{incumb}]}_{\text{1. Skill accumulation of } t-1 \text{ stayers}} \\
 & + \underbrace{\frac{p_{k,t-1}^{lvr} + p_{k,t}^{ent}}{2} \cdot \left(\mathbb{E}[s_{k,i,t}^{ent}] - \mathbb{E}[s_{k,i,t-1}^{lvr}]\right)}_{\text{2. Churning: difference entrants in } t, \text{ leavers after } t-1} \\
 & + \underbrace{\left(p_{k,t}^{ent} - p_{k,t-1}^{lvr}\right) \cdot \left(\frac{\mathbb{E}[s_{k,i,t}^{ent}] - \mathbb{E}[s_{k,i,t}^{incumb}]}{2} + \frac{\mathbb{E}[s_{k,i,t-1}^{lvr}] - \mathbb{E}[s_{k,i,t-1}^{sty}]}{2}\right)}_{\text{3. Marginal selection}}.
 \end{aligned}$$

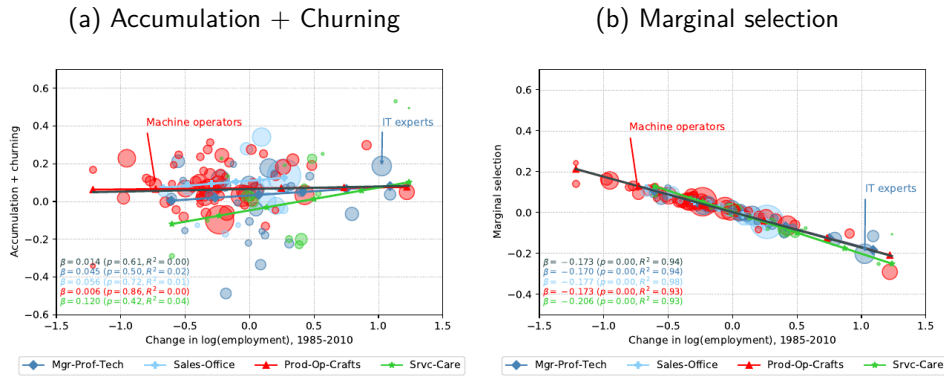
See Section 1.E.1 of the Online Appendix for the steps of the derivation. The first term of Equation (1.13) reflects the skill accumulation of workers who remain in the occupation. Its impact on occupational skill changes is high if turnover is small and skill accumulation of staying workers is high.²¹ The second term is churning, which is composed of average turnover multiplied with the skill differences between period- t entrants and $t-1$ leavers. This will tend to be negative since leavers will have accumulated some skills relative to entrants. It becomes more negative for high turnover occupations and for large estimates of $\Gamma_{k,k}$. Hence, the accumulation and churning effects will often act in opposite directions.²² Importantly, occupation growth does not have a first order effect on either accumulation or churning. By inducing variation in turnover $\frac{p_{k,t-1}^{lvr} + p_{k,t}^{ent}}{2}$, differences between the numbers of entrants and leavers push terms 1. and 2. in opposite directions.

In contrast, occupation growth directly enters the marginal selection effect in the third term of Equation (1.13), which is the product of net entry and the difference in skills between marginal and inframarginal workers in an occupation. In fact, since we have documented above in Section 1.2.2 that, in all occupations, entrants' wages are lower than incumbents' wages and leavers' wages are

²¹In our setup, it will be generated from the estimated $\hat{\Gamma}_{k,k}$ -coefficients and worker demographics.

²²If an occupation is stable in the sense that employment size and skill composition are constant, they must cancel each other out. The marginal selection effect will be zero because of constant employment and the left-hand-side of (1.13) will be zero because of constant skills. Skill accumulation of staying workers must equal the churning effect due to the difference in skills between entrants and leavers.

Figure 1.8: Employment growth vs. the components of skill changes



Notes: Results correspond to sample averages following Equation (1.13). The horizontal axis in both panels shows the change of the log number of employed workers within an occupation between 1985 and 2010. One bubble represents one of the 120 detailed occupations in the SIAB SUF. The four groups show an aggregation of these detailed occupations as described in Appendix Table 1.4. Bubble size corresponds to the number of employed workers in an occupation averaged across years 1985 until 2010. Regression lines across all occupations (black) and within the four broad groups (colored) are weighted by the number of employed workers.

lower than stayers' wages, occupation growth will determine its sign.²³ The marginal selection effect is negative for growing occupations; it is positive for shrinking occupations; and it is zero when there is no change in size. Marginal selection thus formalizes and quantifies the intuition developed in Section 1.2, whereby the more an occupation grows, the more net entry of less skilled workers it experiences.

Figure 1.8 plots employment growth against the sum of the accumulation and churning effects (Panel a) and against the marginal selection effect (Panel b). The patterns are strikingly different. Accumulation and churning are much more dispersed and there is no systematic relation with employment growth. The average is significantly above zero. Since the accumulation effect is generally positive and the churning effect generally negative (see Section 1.E.2 of the Appendix for separate plots), this means that the accumulation effect dominates overall. This is not surprising given that the German working age population grew significantly older and more experienced over the period under study.

Figure 1.8b displays a substantively different pattern. There is a strong negative relationship between employment growth and marginal selection. All 120 occupations are very close to the overall regression line and the four separate regressions are almost on top of each other. This implies that there is no large

²³Note that skill prices are the same for entrants/incumbents in t and for stayers/leavers in $t - 1$. Furthermore, both summands in the second term of the marginal selection effect are negative. Hence, knowing wages is enough to determine the sign of this second term; any particular estimate of skills only affects its magnitude.

variation across occupations in the second term of the product in 3., i.e., the sum of the differences between entrants/entrants/incumbents and stayers/leavers. However, the absolute level of this term is large and induces strong differences in marginal selection between growing versus shrinking occupations. Considering how Figure 1.7b is related to its components in Figure 1.8, marginal selection in the right panel determines the entire negative slope between skill changes and employment changes. The location of the regression line and the variation around it stem from accumulation and churning in the left panel. Therefore, the systematic part of the large selection effects we found in Section 1.4.1 can be traced back to changes in occupations' sizes multiplied with the (negative) skill differences between marginal and inframarginal workers.²⁴

The marginal selection effect lends itself to further analysis. We can split it up into the contributions at entry on the one hand and when leaving the occupation on the other hand. Doing so reveals that the slopes of the regression lines in Figure 1.8b are made up of steeper slopes for entrants versus incumbents and flatter slopes for leavers versus stayers in the three growing occupation groups; they are the same for the shrinking Prod-Op-Crafts (see Appendix 1.E.2). This might not be surprising given that skills should be lower at occupation entry and that leavers should have a larger impact in shrinking occupations. Digging deeper into this, Tables 1.1 and 1.2 decompose the marginal selection effects in two different ways, using the four broad groups for ease of exposition. See Appendix 1.E.3 for the decomposition formulas.

Table 1.1 breaks down the contributions to marginal selection for the broad occupation groups by the origin or destination of marginal workers. Maybe not surprisingly, the single largest contributors are labor market entrants, who make up at least 35% of the total for the three growing occupations and almost one fourth for Prod-Op-Crafts. The main reason for this is that new labor market entrants are a substantial share of entrants into growing occupations and that they have not accumulated much skills in the respective occupation yet. In general, there are some striking differences between the three growing occupation groups on the one hand and Prod-Op-Crafts on the other hand. For the latter, switches to and from unemployment are particularly important. They make up almost 50%; the effect on average skills is positive because of net outflows. Entrants from unemployment account for the same share of marginal selection in Prod-Op-Crafts as sample entrants. In contrast, the combined contribution of unemployment is around 20% for the other three occupation groups and it negatively affects average skills because of net inflows. Leavers to outside of the labor force have a fairly large effect everywhere, that is, a substantial amount of

²⁴In Hsieh et al. (2019), increasing wages per efficiency unit of skill in an occupation also attract workers of lower quality. Their modeling setup equates overall selection effects with marginal selection. With Fréchet as a specific multidimensional skill distribution, Hsieh et al.'s setup implies that selection just offsets the increasing wages per efficiency unit. Interestingly, this implication is approximately borne out in our approach, which does not make a distributional assumption.

Table 1.1: Contributions to marginal selection by origin and destination activities

		Mgr- Prof- Tech	Sales- Office	Prod- Op- Crafts	Srv- Care
Entrants	Switchers from Mgr-Prof-Tech		-0.05	-0.00	-0.00
	Switchers from Sales-Office	0.01		0.01	0.01
	Switchers from Prod-Op-Crafts	0.13	0.11		0.07
	Switchers from Srv-Care	0.01	0.02	0.02	
	From unemployment	0.11	0.10	0.22	0.11
	From outside of the labor force	-0.01	-0.02	0.09	-0.04
	Sample entrants	0.38	0.48	0.23	0.35
Leavers	Switchers to Mgr-Prof-Tech		-0.05	-0.01	0.00
	Switchers to Sales-Office	0.05		0.01	0.02
	Switchers to Prod-Op-Crafts	0.10	0.14		0.15
	Switchers to Srv-Care	0.02	0.03	0.01	
	To unemployment	0.07	0.07	0.26	0.11
	To outside of the labor force	0.20	0.21	0.15	0.23
	Sample leavers after age 54	-0.06	-0.04	-0.01	0.01

Notes: The numbers represent relative contributions to the marginal selection effect within each broad occupation group during 1985–2010. Columns sum to one. The actual sizes of the effects are -0.03 (Mgr-Prof-Tech), -0.02 (Sales-Office), 0.06 (Prod-Op-Crafts), and -0.04 (Srv-Care). The explicit decomposition formulas are in Appendix 1.E.3.

less-skilled workers are leaving all occupation groups in each period. However, entrants from outside the labor force exert a counteracting effect on marginal selection for the growing occupation groups; they often enter from other forms of employment that are not covered in our data (self-employment, civil servants, work abroad) and they are quite high-skilled compared to incumbents. Leavers after age 54 also mostly exert a counteracting effect as they have accumulated substantial skills over their careers and they exit our sample for exogenous reasons.

Quantitatively, most of the marginal selection effect in Table 1.1 is accounted for by moves into or out of unemployment, the labor market, or the sample.²⁵ Nonetheless, direct switches between occupations are non-negligible and of economic interest because they almost always positively contribute to marginal selection. That is, entrants into an occupation, independent of the

²⁵The lion's share of these moves is part of transitioning between jobs. In Online Appendix 1.F.1 we repeat Table 1.1 for our sample where we have filled non-employment spells using the wage and occupation of the adjacent spell with the lower wage. We find that the role of switches between occupations approximately doubles (there are still permanent entry and exit from the sample as alternative contributors). We will come back to this and the robustness of our results in Section 1.4.3.

origin occupation, are less skilled than the incumbents. Leavers from an occupation, independent of the destination occupation, are less skilled than the stayers. The partial exception is Mgr-Prof-Tech, where switchers from that occupation group are more skilled than the incumbents in their destination and switchers into Mgr-Prof-Tech tend to be more skilled than the stayers in the respective origin occupations.²⁶ But overall this evidence once again indicates that, at the time of switching, incumbents and stayers have strong specific skills in their occupation compared to marginal workers, which is hard to reconcile with a one-dimensional ranking of occupations by skill. Online Appendix 1.B.2 rejects the one-dimensional skill model in our data based on this type of evidence.

Table 1.2 shows the contributions to marginal selection by the sources of workers' skills. We employ the longitudinal information in the data to separate workers' skill endowment at the most recent entry into the occupation from their skill accumulation since then. In particular, we calculate the endowment from observed wages and normalized prices at the time of entry. We then obtain predicted skill accumulation during the current stint by summing the respective estimated $\hat{\Gamma}_{k,k}$ over the worker's tenure. Finally, we calculate the deviation of workers' actual wages from our prediction based on systematic skill and price changes. The first row in Table 1.2's top panel shows that in all occupation groups, entrants have lower skill endowments than incumbents had at the time that they were entrants. The corresponding bottom panel shows that also leavers are negatively selected relative to stayers when comparing endowments retrospectively.²⁷ The resulting contribution to marginal selection that is due to different skill endowments is substantial, ranging from one fifth in Mgr-Prof-Tech and Sales-Office to two thirds in Srvc-Care. Endowments at entry into the occupation can be interpreted as a "classic selection effect", i.e., as in cross-sectional models where workers' skill endowments are drawn before making the occupational choice.

Workers' skills do however change during their stint in an occupation and this has important separate effects on marginal selection. Table 1.2 shows that for the high-accumulation Mgr-Prof-Tech and Sales-Office occupations, around 40% of marginal selection is due to skills accumulated by incumbents (accumulation for entrants is zero by construction).²⁸ Another 15–20% percent is due to more skills accumulated by stayers compared to leavers (i.e., stayers' tenure is on average longer than that of leavers). Not surprisingly, the magnitude

²⁶The latter makes sense as, e.g., promotions to team leader might yield such a change of occupation.

²⁷Comparing endowments in Table 1.2 presents a model-consistent way to control for experience in Figure 1.3 above. Also by this measure, both entrants or leavers earn less than incumbents or stayers.

²⁸Notice that, while the skill differences due to skill accumulation are in principle temporary, higher values for incumbents will continue to contribute towards the marginal selection effect as long as the respective occupation keeps growing and drawing in new, less skilled workers.

Table 1.2: Contributions to marginal selection by source of skills

		Mgr- Prof- Tech	Sales- Office	Prod- Op- Crafts	Srv- Care
Entrants	Endowment at the most recent entry into the occupation group	0.14	0.16	0.21	0.33
	Predicted skill accumulation since the most recent entry	0.40	0.41	0.24	0.08
	Deviation of skills from the prediction since the most recent entry	0.10	0.07	0.13	0.08
Leavers	Endowment at the most recent entry into the occupation group	0.08	0.06	0.18	0.35
	Predicted skill accumulation since the most recent entry	0.16	0.19	0.04	0.09
	Deviation of skills from the prediction since the most recent entry	0.13	0.11	0.20	0.08

Notes: The numbers represent relative contributions to the marginal selection effect within each broad occupation group during 1985–2010. Columns sum to one. The actual sizes of the effects are -0.03 (Mgr-Prof-Tech), -0.02 (Sales-Office), 0.06 (Prod-Op-Crafts), and -0.04 (Srv-Care). The explicit decomposition formulas are in Appendix 1.E.3.

due to skill accumulation is very heterogeneous and it accounts for less than 30% of the marginal selection effect in Prod-Op-Crafts and for even less in Srv-Care. This heterogeneity underlines the importance of flexibly modeling skill accumulation across occupations: In a model with homogeneous returns to experience the effect would only depend on tenure and entry age, which vary much less across occupations and may even lead to inverse predictions such as overall accumulation being highest in Prod-Op-Crafts.²⁹

The final contributor to marginal selection in Table 1.2 are deviations from what is already captured in our estimates $\hat{\Gamma}_{k,k}$. These deviations are due to systematically different skill shocks of incumbents/stayers compared to entrants/leavers during the stint.³⁰ Again, the value is zero by construction for entrants, showing that incumbents are positively selected on this margin as well. The same holds true for stayers versus leavers. The effects are quantita-

²⁹In our data, Prod-Op-Crafts occupations exhibit the longest average tenure (more than 14 years), Srv-Care the shortest (10 years), and the other two are right in the middle (just below 12 years). Abstracting from differences in entry age, the skill accumulation effect on marginal selection would be highest for Prod-Op-Crafts. In contrast, the longer average tenure does not translate into a quantitatively large effect in our model because the skill accumulation coefficients are much lower for Prod-Op-Crafts than for Mgr-Prof-Tech or Sales-Office (see Figure 1.6).

³⁰As discussed in Section 1.3.4, differential employer learning about skills is an alternative explanation.

tively substantial and economically interesting. They show that the argument we made in Section 1.3—that staying in an occupation is endogenous in the sense that only workers who receive sufficiently favorable skill shocks will decide to remain in it—is not merely an academic one. The deviations are consistent with learning models of occupational mobility, such as Groes et al. (2014) and Papageorgiou (2014), which show that workers who leave an occupation previously systematically deviate from their peers in terms of wages (expected skills). While predicted skill accumulation is quantitatively large, stayers in occupations are clearly selected according to their idiosyncratic skill shocks. This underscores the importance of the self-selection model underlying our estimation method.³¹

To sum up the evidence from this section, marginal selection can account for the systematic part of the relationship between skill changes and employment changes across occupations implied by our estimates, both qualitatively and quantitatively. This selection effect fundamentally stems from the fact that entrants and leavers are substantially less skilled than incumbents or stayers in any occupation; it is *due to* sector growth. Marginal selection also conforms with reasonable economic notions of the underlying selection effects. First, a large part of it are moves into and out of employment. Second, most almost all groups of switchers contribute negatively (positively) to skills changes in the destination (origin) occupation. The skill differences of marginal versus infra-marginal workers consist of differing endowments at entry, skill accumulation of incumbents, and endogenous switching (staying) of those workers who experience negative (positive) shocks during their stint in an occupation. These effects are strong and seem economically plausible. As their magnitude does not depend much on our skill price estimates, we consider the results in this section separate and substantive evidence in favor of the results from our estimation method.

1.4.3 Robustness of Results and the Task Content of Occupations

Appendix 1.F examines the robustness of our empirical results in alternative samples and estimation specifications. We briefly summarize the reasons for and results of these robustness checks in the following. Finally, we connect to prior literature by describing occupations via the task content of work.

³¹This is a case where the more general acceleration or deceleration interpretation of skill price changes from Section 1.3.3 matters. In particular, if skill prices had already risen in the base period, we would overestimate the difference in skill accumulation of entrants versus incumbents (and leavers versus stayers), while underestimating the difference in endowments at the most recent entry. The (economically instructive) role of deviations from the model prediction is unaffected by this more general interpretation, however. See Online Appendix 1.E.3 for more details.

Filling non-employment spells: In our view, a key robustness check is to allow for endogenous unemployment and exit from the labor force. For instance, when the skill price in an occupation declines, workers might prefer to temporarily leave employment over switching to another occupation directly if the benefits they obtain in unemployment are sufficiently high. As an alternative, we therefore include all intermittent non-employment spells in our sample by imputing workers' wages and their occupation choices. We do this by comparing wages before and after the non-employment spell, and assign workers the lower of those two wages adjusted for inflation as well as the corresponding occupation. That is, we assume that workers could well have worked in the lower paying occupation but decided to become unemployed or exit the labor force for some time instead.

Re-running the entire analysis on the sample constructed in this way, we find that the estimated skill accumulation coefficients are generally smaller and they turn negative in cases that one would expect to be "downward" switches (e.g., from Mgr-Prof-Tech to Prod-Op-Crafts). Yet the other results are similar to before: The correlation between wage and employment growth is approximately zero but it is strongly positive between price and employment growth (though slightly flatter than in the main sample). In addition, the implied skill changes are again negative and closely related to marginal selection. Section 1.5 returns to this filled sample for some of the wage inequality analyses.

Different demographic groups: We have restricted our main sample to prime age West German men as these can be defined consistently over the 1975–2010 period and many potentially confounding factors (e.g., rising participation and education rates, changing discrimination) do not apply. Nonetheless, it is still informative to see whether the broad results hold up when we change the demographics. First, in a wide definition that adds women, East Germans, and workers who are always foreigners, the results are very similar to our main sample.³² If we consider West German women only, it is striking that the employment distribution is very different, with substantially more Sales-Office and Srvc-Care occupations (indicated by the bubbles sizes in the respective graphs). Nonetheless, the results for the women sample are similar to our main results: there is no relationship of employment with wage growth but with skill price growth (even slightly stronger), while implied skill changes and marginal selection again point in the same direction. Finally, restricting ourselves once more to West German men, but extending the age range to 20–60, our original findings are confirmed with somewhat steeper slopes. This makes sense as very young workers had less time to accumulate skills on their jobs and early retirement—which was important over many of the years in our sample period—was likely to be selective.

³²The same holds if we conversely exclude anyone who was ever coded as a foreigner; naturalized citizens who change their foreigner status are discussed further in Section 1.5.

Different estimation specifications: We also estimate different model specifications that were discussed in Section 1.3. First, we employ the identification approach pioneered by Heckman, Lochner, et al. (1998), which assumes that mature workers' skill growth should be rather flat. Our estimated coefficients depicted in Figure 1.6 lend support to this assumption. At the same time, possibly forward-looking choices—i.e., via picking occupations with high skill accumulation—should be less of a concern in this age group. We thus follow Heckman, Lochner, et al. (1998)'s flat spot approach and set skill accumulation to zero for the 45–54 year old subsample. This obtains a similarly steep relationship between skill price and employment changes as in the baseline estimation. Marginal selection also works in the same direction as the implied skill changes, although it is somewhat flatter.

We then return to our main sample and enrich the skill accumulation function to be education-group specific. The results hardly change compared to the baseline where occupational skill accumulation differs only across age-groups. Next, we allow for changing non-pecuniary amenities in occupations by augmenting the estimation equation with regressors for occupation switches (detailed derivation in Appendix 1.B.3). We perform this exercise for the four broad occupations only because of the extensive data requirements. Similar to Hsieh et al. (2019), we find that changing amenities (or, alternatively, changing preferences) hardly have an effect on the estimated changes of skill prices.³³ We also estimate the alternative occupation-specific fixed effect approach by Cortes (2016). The relationship between employment and price changes is somewhat flatter than in our proposed method but still highly significant.

Notably, in each of the different samples and estimation specifications the relationship between occupational wage and employment growth is essentially flat, whereas estimated skill prices and employment growth correlate positively. This indicates that demand shifts were dominant across broader demographic groups and that selection effects are generally strong, masking this underlying driving force. Moreover, we find that the estimated price changes positively correlate across samples and estimation specifications, and that implied skills and marginal selection work in the same direction throughout.

Connecting to the task-based approach Finally, we connect to a large literature that has investigated occupational changes with the task-based approach (e.g., Autor, Levy, et al., 2003; Firpo et al., 2013; Yamaguchi, 2018). For Germany, several authors have used the Qualifications and Career Surveys (QCS)

³³Hsieh et al. (2019) introduce a general equilibrium Roy model in which workers sort on either (unobserved) talent or preferences. They report a small, weakly positive correlation between occupational employment and earnings growth similar to Figure 1.2a above (Figure 10, p. 41 Hsieh et al., 2019). Based on that finding in combination with their model prediction, they conclude that workers primarily sort into occupations based on talent as opposed to preferences.

to measure routine versus non-routine task content in particular (e.g., Spitz-Oener, 2006; Antonczyk et al., 2009; Gathmann and Schönberg, 2010). In the last section of Appendix 1.F, we also employ the QCS to construct routine as well as analytical, interactive, and manual task content for our 120 occupations. We then relate these measures to employment growth, wage changes, and our estimates for prices and skills.

The resulting graphs show that occupations intensive in analytical (often in the Mgr-Prof-Tech group) and interactive (Mgr-Prof-Tech and Sales-Office) tasks indeed grew quite strongly, whereas employment in routine-intensive (Prod-Op-Crafts) occupations declined. High analytical and interactive task content of occupations helps predict rising wages. However, the relation with estimated skill prices is even steeper. Conversely, implied skills deteriorate in analytical and interactive task content. The correlation between routine task intensity and average wages is zero; this is composed of falling prices and rising skills. All this is consistent with the impact of RBTC on these occupations and with our finding that skill price changes are counteracted by selection effects.

The case of manual-task intensive occupations (mostly in the Prod-Op-Crafts and Srvc-Care groups) is also in line with the latter general finding. But it seems that the overall demand shift was negative because employment as well as average wages and skill prices declined. One likely reason for this is measurement, since the QCS questionnaires have some difficulty distinguishing between routine and manual job tasks. The other is that alternative demand forces than RBTC have lifted the employment and skill prices of Srvc-Care occupations, despite their high (measure of) manual tasks.³⁴

These results demonstrate the usefulness of task measures as a dimension-reduction device, particularly when working with more limited datasets. It is especially helpful to study specific drivers of occupational change. At the same time, using detailed occupations directly is most flexible and does not require precise measurements of all task dimensions for which demand may have changed.³⁵

1.5 Skill Prices and Wage Inequality

We have shown that selection effects largely explain why occupational wages and employment growth are uncorrelated over the period under study. By a similar token, selection may shroud the relation between demand shifts and wage

³⁴Additional forces that could have worked on Srvc-Care include demand for social skills or consumption of low-skill services (Autor and Dorn, 2013; Mazzolari and Ragusa, 2013; Deming, 2017). In the case of Prod-Op-Crafts occupations, employment may have declined even more than predicted by RBTC because of trade and offshoring (Autor, Dorn, and Hanson, 2013; Goos, Manning, and Salomons, 2014). See also Footnote 2.

³⁵For example, Firpo et al. (2013), Blinder and Krueger (2013), and Goos, Manning, and Salomons (2014) construct task measures for offshorability in the U.S. and Europe. Deming (2017) constructs measures of social skills.

inequality, particularly between occupations. In this section, we thus examine to what extent selection may also be responsible for the result that occupations exhibit limited explanatory power for the increase of wage inequality. Running Mincer-regressions, Card et al. (2013) obtain only small decreases in residual wage inequality when adding occupation dummies. Dustmann, Ludsteck, et al. (2009) find that in the lower half of the wage distribution, occupational demand is not a first-order factor driving wage differences. We first use only the estimated skill prices and selection to quantify the forces driving between-occupation inequality. We then employ the full version of our model to disentangle the components that affect various percentiles of the wage distribution,³⁶ paying particular attention to entry wages, demographics, skill accumulation, and prices.

1.5.1 The Attenuating Effect of Selection On Inequality

Over the period of our study, the variance of log wages multiplied with 100 went up by 12.4 points from a baseline of 14.3. The component due to differences between occupations started at a value of 5 in 1985. It then more than doubled and reached almost 40% of the overall inequality in 2010. A substantial share of the increase thus occurred between occupations, consistent with occupational demand (e.g., due to routine-biased technical change and offshoring as in Acemoglu and Autor, 2011) but also with other factors having been important drivers of wage inequality.

In particular, one question that several papers before us have asked is whether changes in the demographic structure of the population were such a factor. The first column of Table 1.3 reports on a counterfactual analysis similar in spirit to that of Figure 16 in Autor (2019).³⁷ Holding wages at their 1985 level, we reweight observations with the distribution of age, foreigner status, and education in 2010.³⁸ This exercise answers the question: what if choices conditional on these observables and wages were constant at their 1985 levels, but the demographic structure had shifted to that of 2010 (due to population aging, increased immigration, and rising educational achievements of younger

³⁶Due to the nature of our data, we restrict attention to wage inequality as opposed to overall inequality. It is thus important to note that we do not see a clear trend in labor force participation rates of German men over most of the period under study. In particular, there was a decline between 1975 and 1989, but rates stabilized around 93-94% thereafter. This is in stark contrast to U.S. men, where rates dropped from almost 94% to below 90% between 1989 and 2010.

³⁷The closest to his analysis includes occupational choices among the variables used for reweighting. The results of this exercise can be found in Appendix 1.G.1. They are quantitatively similar to our specification in Table 1.3. We prefer this specification because age, foreign, and education are arguably all factors that mostly contribute to occupational supply as opposed to demand.

³⁸To compute the weights, we follow DiNardo et al. (1996) using a logit model with 30 dummies for detailed ages between 25–54, a dummy for being permanently German or not, as well as three dummies for education status.

cohorts)? Quantitatively, the answer is similar to that of Autor (2019) in that the effects make up something like a fifth of the total increase and a third of the increase in between-inequality.

Table 1.3: Decomposition of the between-variance of wages, data and counterfactuals

	Counterfactuals			Actual
	Rewgt. age, foreign, educ.	Prices only	Prices + rewgt. age, foreign, educ.	Data + price estimates
Overall $\Delta\sigma^2(w_{i,t})$	2.41	5.13	9.56	12.41
Between $\Delta\sigma^2(\bar{w}_{k,t})$	1.74	5.13	8.88	5.25
$2 \cdot \sigma(\Delta\bar{w}_{k,t}, \bar{w}_{k,1985})$	$2 \cdot \sigma(\Delta\pi_{k,t}, \bar{w}_{k,1985})$	0.00	3.23	3.23
	$2 \cdot \sigma(\Delta\bar{s}_{k,t}, \bar{w}_{k,1985})$	0.53	0.00	-0.53
$\sigma^2(\Delta\bar{w}_{k,t})$	$\sigma^2(\Delta\pi_{k,t})$	0.00	1.89	1.76
	$\sigma^2(\Delta\bar{s}_{k,t})$	1.21	0.00	3.02
	$2 \cdot \sigma(\Delta\pi_{k,t}, \Delta\bar{s}_{k,t})$	0.00	0.00	-2.24

Notes: All values are multiplied with 100. The levels in 1985 are 14.3 (overall) and 5.0 (between). Based on specification with 120 occupations. $\bar{w}_{k,t}$ refers to the average wage in occupation k in year t . The counterfactual experiments are: *Rewgt. age, foreign, educ.*: take observations in 1985 and reweight them to match the 2010 distribution of these characteristics with weights computed following DiNardo et al. (1996) in order to obtain 2010 wages. *Prices only*: take individual wages in 1985 and add our estimated price changes to obtain 2010 wages. *Rewgt. age, foreign, educ. + prices*: Combine both experiments.

With the help of our model we can gain further insights into the components that have driven changes of between inequality. Denoting average wages (skills) in an occupation by \bar{w}_k (\bar{s}_k), we can write the change between 1985 and 2010 as:

$$\begin{aligned} \Delta\sigma^2(\bar{w}_{k,t}) = & \underbrace{2 \cdot \sigma(\Delta\pi_{k,t}, \bar{w}_{k,1985}) + 2 \cdot \sigma(\Delta\bar{s}_{k,t}, \bar{w}_{k,1985})}_{2 \cdot \sigma(\Delta\bar{w}_{k,t}, \bar{w}_{k,1985})} \\ & + \underbrace{\sigma^2(\Delta\pi_{k,t}) + \sigma^2(\Delta\bar{s}_{k,t}) + 2 \cdot \sigma(\Delta\pi_{k,t}, \Delta\bar{s}_{k,t})}_{\sigma^2(\Delta\bar{w}_{k,t})} \end{aligned} \quad (1.14)$$

See Appendix 1.G.1 for the detailed derivation. First consider the terms underneath the braces, which involve only wages. They say that if overall wage inequality was constant $\Delta\sigma^2(\bar{w}_{k,t}) = 0$ and there were any changes in the wage structure across occupations $\sigma^2(\Delta\bar{w}_{k,t}) > 0$, it must be that occupations at the bottom of the distribution experienced better wage growth than those at the top on average $\sigma(\Delta\bar{w}_{k,t}, \bar{w}_{k,1985}) < 0$. Using our model, the main terms in (1.14) now break these two components into changes of prices and changes of skills. We start by applying this decomposition to various counterfactual experiments in order to better understand the mechanisms at work.

The remainder of the first column of Table 1.3 shows that the reweighting procedure only affects skills; all terms involving price changes are zero. The covariance of skill changes with baseline wage levels is positive. This is a reflection of the fact that the population grew older and more educated together with high-wage, high-education occupations (Mgr-Prof-Tech, Sales-Office) featuring faster skill accumulation over the life cycle. However, while we will show below that these changes in the demographic structure had an important role for overall inequality, they play a limited role for explaining between-occupation inequality. Including occupations among the variables used for reweighting does not change these conclusions (see Appendix 1.G.1).

The second column of Table 1.3 reports on the results from the opposite experiment, which isolates the effect of price changes. Holding constant the 1985 demographic structure and occupation choices, we add the cumulative changes of occupational skill prices between 2010 and 1985 to individuals' wages. These effects alone generate almost the entire increase of between inequality. The bulk of the effect stems from the covariance between price changes and initial wage levels. Prices rose in Mgr-Prof-Tech, Sales-Office, and Srvc-Care occupations; the first two featured high wages already in 1985 and employment there is much larger than in Srvc-Care. Our preferred interpretation of this term is that it reflects the nature of demand shifts: during the period under study, they happened to benefit high-wage occupations more. Naturally, any term involving skill changes is zero.

The third column shows what happens if both experiments are turned on. The variance of price changes rises somewhat due to the different weights; all other effects from the separate experiments remain the same. The covariance between price and skill changes is substantial and positive. Overall, this counterfactual overestimates the rise of between inequality by two thirds. Looking at the first line only, one may even be tempted to think that this exercise explains a large share of the overall rise in inequality (77%).

However, a comparison with the last column, the actual between variance and its components, reveals that this large "explained" share is far off, since there are important dampening effects of selection on wage inequality. In particular, the economic mechanism described at length in the previous section—a deterioration of skills in occupations where prices rose—has a strong impact on inequality. This strong negative covariance is everything but mechanical: if we interpret price changes as mainly driven by demand shifts and the demographic changes captured by the reweighting as supply shifts to occupations, the third column of Table 1.3 suggests that these shifts covaried positively. The impact of -2.24 points in the actual data as well as the negative covariance of skill changes and initial wage levels are therefore important attenuating selection effects. As a result, the actual contribution of skill changes to between inequality is negligible whereas in the counterfactual it is $+3.29$ points overall.

What we have just described is consistent with theoretical results by Heckman and Honoré (1990) for a two-sector Roy economy. They showed that,

if the population distribution of skills is log concave, self-selection in the Roy model will generally lead to more equal wages compared to random assignment into occupations. With respect to the particular case at hand, when the correlation of skills in the different occupations is sufficiently low, average skills in the occupation with declining prices will unambiguously improve and they will unambiguously deteriorate in the occupation with rising prices.

Our results show why decompositions based on observables alone have difficulties generating quantitatively meaningful increases of inequality: so long as average wages across occupations are more or less constant (e.g., see again Figure 1.2), changing demographics and even large shifts of the employment structure across occupations exert limited impact. The reason is that underlying skill prices and supply changes, which would have raised between-occupation inequality further than what is observed, are counteracted by strong selection effects.

1.5.2 Factors Contributing to Wage Inequality

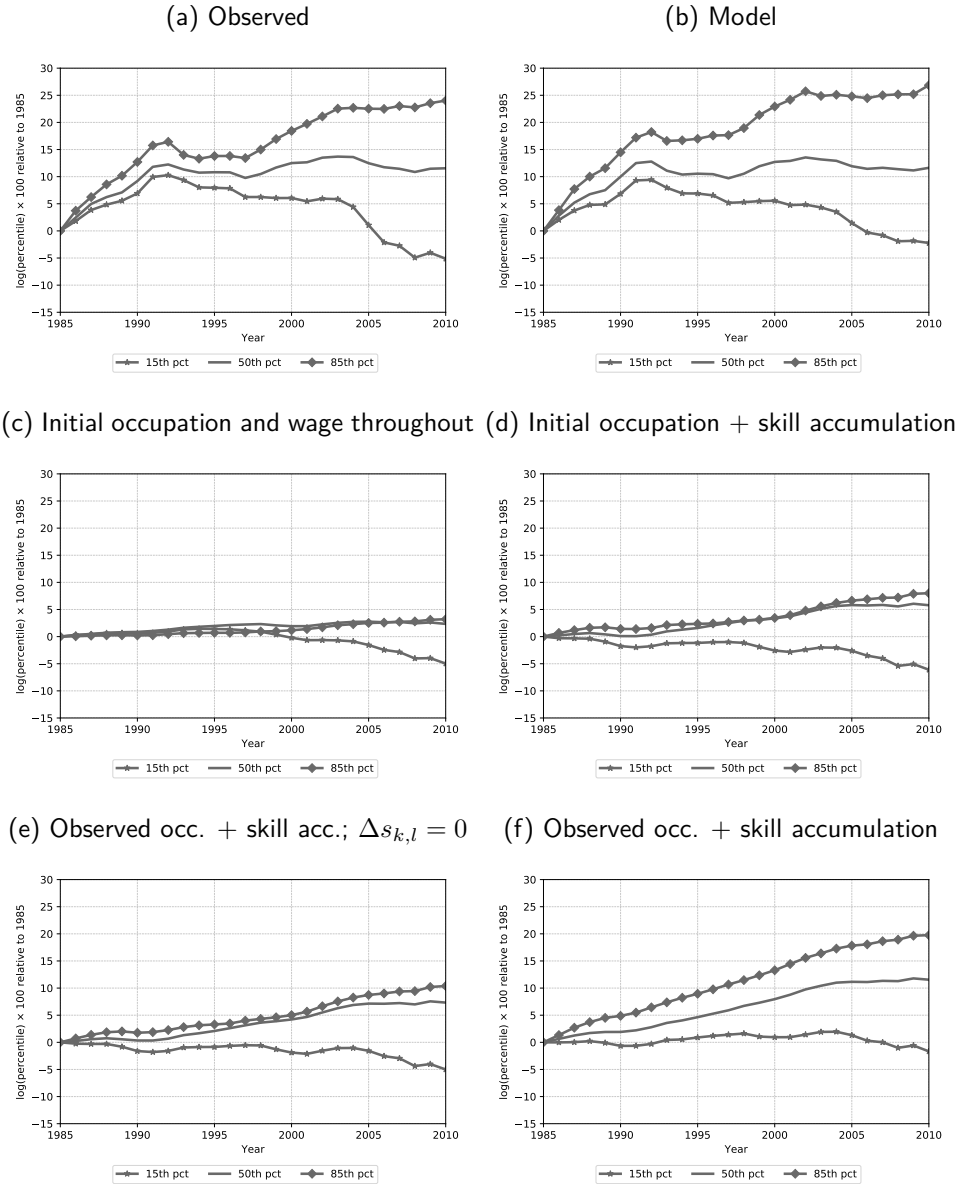
While the economic forces under scrutiny in this paper are most important for inequality between occupations, our model can be employed to gain a better understanding of the overall development of the wage distribution, too. In the following, we use our estimates to disentangle the factors that contributed to differences between the quantiles of the wage distribution.

Figure 1.9 plots the evolution of the percentiles of the wage distribution in the data and in various scenarios based on our model. Figure 1.9a just repeats Figure 1.1a for ease of comparison; it shows the strong widening of the German wage distribution (Dustmann, Ludsteck, et al., 2009; Card et al., 2013). Figure 1.9b plots the individual-level predictions from our model. In order to obtain an individual's predicted wage in a particular year, we start from the initial wage observed in our data³⁹ and follow his occupational choices over the life-cycle, adding the relevant skill accumulation parameters and skill price estimates along the way. The predictions track the data closely, both qualitatively and quantitatively. Note that all percentiles in all panels are normalized to zero in 1985; Table 1.25 in the Appendix shows that our model is close to its targets also for the levels of these percentiles and the variance.

The remaining panels of Figure 1.9 investigate the drivers of the model prediction by starting with the most basic version and turning on our model's features one after the other. Panel c reports on how the three percentiles would have evolved if workers had kept their initial wages for their entire working life. Many variants of supply changes would directly affect this scenario. For example, one may expect the expansion of tertiary education to lead to higher entry wages

³⁹For workers who may have entered the labor market before our sampling period starts—i.e., those born before 1950 observed to be working in 1975—we use our skill accumulation estimates to impute their initial wages at age 25, assuming they stayed in the same occupation all along.

Figure 1.9: Wage inequality scenarios



Notes: Panel a: observed wages. Panel b: simulated life-cycle trajectories based on our full model: starting from the initial wage and occupational choice, add all skill accumulation and price change estimates using occupational choices observed in the data. Panel c: workers keep their initial wage throughout the life cycle. Panel d: workers stay in their initial job throughout the life-cycle; in each period, we add the skills they would have accumulated in that job (i.e., $\Delta s_{k_0, k_0, t}$). Panel e: use observed switches, setting direct gains from switching to zero, i.e., $\Delta s_{k,l} = 0 \forall k \neq l$. Price changes are zero as well, so the difference to Panel d comes purely from differential skill accumulation in occupations. Panel f: as in Panel e, but adding the direct gains from switching. The only difference to the full model in Panel b are the price changes, which continue to be zero. In all scenarios, we treat unemployment or out-of-the-labor force spells as follows: when such a spell is observed in the data, simulated workers do not enter the inequality statistics. Furthermore, we assume no depreciation and upon re-entry into paid work add—where relevant—the $\Delta s_{k,l,t}$ with l being the occupation before the spell.

for the additional university graduates, raising the upper percentiles. The results show that the median and the 85th percentile rose somewhat. Quantitatively, this is not very important, making up between one fifth (median) and one eighth (85th percentile) of the total increase. All three percentiles evolve rather smoothly, the distinct temporal pattern over time visible in Panels a and b thus does not seem to be driven by changing conditions at labor market entry.

After a small initial increase, the 15th percentile exhibits a pronounced decline starting in the mid-nineties. In fact, this decline is so strong that it could explain the drop of that percentile between 2010 and 1985.⁴⁰ This large drop seems due to temporary workers⁴¹ and naturalized citizens, both of whom are frequently the same. Excluding workers ever coded as foreigners from our sample reduces the fifteenth percentile drop by more than two thirds both in Panels b and c (see Figure 1.53 in the Appendix). This is consistent with Dustmann, Ludsteck, et al. (2009)'s hypothesis that, from the 1990s onward, many low-skilled immigrants and ethnic Germans from Eastern Europe increasingly flowed into the West German labor market, worsening the composition of employment at the lower end.

In Figure 1.9d, we continue to assign workers to their initial occupation, now adding the skill accumulation coefficients. There is hardly any change for the fifteenth percentile compared to Panel c, but the median and 85th percentile rise strongly. The incremental changes are 4 points at the median and 5 points at the 85th percentile, amounting to one half (one third) of the overall changes between 2010 and 1985. Again, all changes happen rather smoothly. The scenario shows that the demographic and occupational composition has a quantitatively strong impact on the rise of the upper half of the wage distribution.

We add the empirically observed switches to careers in Figure 1.9e, but do not turn on the direct gains from switching, i.e., we set $\Delta s_{k,l} = 0 \forall k \neq l$. This exercise drives up the median and 85th percentile by an additional three points; it hardly affects the fifteenth percentile. The results show that switches from occupations with flatter age profiles to those with steeper age profiles do matter even if one ignores the oftentimes large jumps associated with switches. However, the skill accumulation differentials are not large enough to drive a majority of earnings inequality. Part of this may be due to timing: For switches that occur after age 35, the skill accumulation differentials between occupation groups are not as big as they are at the beginning of careers. Nevertheless, the rise in the median and 85th percentile is large in Panels d and e in comparison to Panel c. In Appendix 1.G.3, we show that this is due to the aging of the

⁴⁰Again, the temporal pattern is far off. One would reach a radically different conclusion regarding the fit if one were to compare, say, 1991 and 1985.

⁴¹We identify temporary workers from the detailed occupation "assistants without further specification", which mostly appears in the industry group "Credit and insurance intermediation, land and housing, rentals". This industry group contains the subgroup "labor recruitment and provision of personnel" where temporary agencies are listed. Temporary work has increased a lot in Germany (Eichhorst and Tobsch, 2013).

workforce with many more middle-aged workers at the median in 2010 compared to 1985. Demographic factors were therefore substantially responsible for the increase of lower half inequality. In contrast, since the 85th percentile was similarly raised as the median, upper half inequality did not increase much because of demographic changes or skill accumulation within occupations.

Adding the gains associated with occupation changes in Figure 1.9f raises all statistics; it does so disproportionately for the 85th percentile. This is not too surprising given the large coefficient estimates for switches into Mgr-Prof-Tech and Sales-Office occupations that we reported on in Section 1.4. Comparing the end points of our sample period, this scenario explains three quarters of the increase in the 85th percentile and the entire increase in the median; we are too optimistic about the evolution of the fifteenth percentile by 3 points. There are two things to note, however. First, the temporal pattern is very smooth and we do not track the intermittent evolution of any percentile very well. Second, there is actually a decline in the fifteenth percentile for the specification where we make unemployment or exiting the labor force a choice by filling such spells with the lowest adjacent wage (see Figure 1.54 in the Appendix). This suggests that indeed careers at the lower end of the distribution became more fragmented and our main way of treating non-employment spells hides parts of this.

Comparing Figures 1.9f and 1.9b, we see that skill prices explain most of the remaining differences with the actual wage distribution. In particular, changing prices raise the 85th percentile as well as upper half inequality by an additional seven log points. As in the case of between occupation inequality, they thus have a strong impact. In the specification where unemployment is a choice, price changes hurt both the median and the 15th percentile, again highlighting that we overestimate the gains from switching at the lower end because occupation changes involving wage losses often go via an unemployment spell. Finally, adding the price changes allows us to track the temporal evolution of all quantiles. Thus skill prices are not only aligned with employment across occupations, but they also align the temporal patterns of the wage distribution.

In sum, we show that initial occupational choices and demographic factors account for most of the increase in lower half inequality; alternative specifications suggests that more unstable employment biographies and adverse price developments have some role to play, too. This is consistent with the hypothesized effects in Dustmann, Ludsteck, et al. (2009), which is an important finding overall because it has previously been hard to rationalize polarizing demand for occupations together with wage inequality that increased across the board in most countries and time periods (Goos and Manning, 2007; Mishel et al., 2013; Naticchioni et al., 2014; Green and Sand, 2015). Occupational switches and changing skill prices have a particularly important role to play in the upper half of the wage distribution, driving almost all of the additional wedge that opened

up between the 85th percentile and the median over the period 1985–2010.⁴²

1.6 Discussion and Conclusion

This paper develops a model of occupation choice based on Roy (1951), which remains empirically tractable for many occupations and accommodates heterogeneous skill changes over the life cycle. We use this model to study how occupational employment growth relates to occupational wages and overall wage inequality. Our results indicate that skill-constant occupational wages (*skill prices*) evolved in a way that is consistent with occupational demand shifts. Skill selection of workers completely masked this relationship in raw occupational wages, where the development was unrelated to employment changes. We show that the systematic part of the skill price-employment growth nexus is due to what we term the *marginal selection effect*; net entry into an occupation multiplied with the skill differences between occupation entrants/leavers versus incumbents/stayers.

The selection effect that we uncover is more subtle than the one considered in classic Roy models, where workers' skills across occupations are fixed over time. The classic effect accounts for less than forty percent of marginal selection in most occupations. The more important share is due to skill changes during employment stints in an occupation, i.e., the fact that incumbent/staying workers are positively selected, translating into longer tenure and gains from (specific) experience. These effects vary strongly across occupations, which is a reflection of the fact that occupational life-cycle wage profiles are very heterogeneous.

We further show that similar lines of reasoning carry over to wage inequality, where we establish a long-suspected connection to demand shifts and occupational employment changes that is meaningful also in quantitative terms. Worker (self-)selection leads to substantially lower wage inequality between occupations than would be observed if workers in the 1980s were given the skill prices of later decades while holding their occupational choices fixed. Selection thus makes it appear that occupational changes were not that important. Using our model to understand the trends in overall wage inequality, we instead find that differentially evolving skill prices and heterogeneous skill changes across occupations are the most important drivers of upper-half inequality. Initial occupation choice and population aging—which induces higher wages at the median of the wage distribution due to a larger fraction of seasoned workers—are the main factors driving lower-half inequality.

⁴²The discussion in this section is robust to the more general acceleration/deceleration interpretation of skill price changes from Section 1.3.3. First, the full estimated model, which includes both skill accumulation and skill prices, is unaffected by this interpretation. Second, before the estimated prices are included, one would still like to add the average rates of price changes in the base period to the skill accumulation in order to obtain scenarios where “only” entry wages, initial occupations, or occupational switching changed. This is effectively what we do in Panels c–f of Figure 1.9.

Our explanation is consistent with other accounts of rising wage inequality in Germany. One of the most prominent is based on de-unionization and a decentralization of the wage bargaining process (Dustmann, Ludsteck, et al., 2009; Dustmann, Fitzenberger, et al., 2014). These phenomena have the strongest impact in the manufacturing sector, that is, the industry sector that is most important for the declining Prod-Op-Crafts occupations. We deem it plausible that demand shifts are a deeper cause for this, as unions and works councils understand their deteriorating bargaining position due to the threats of substitution by machines or foreign workers.⁴³ Our findings are also consistent with other work showing that German firms tend to upgrade labor through investment in skills (Battisti et al., 2017; Dauth et al., 2017).⁴⁴ These responses may reflect the institutional environment, which results in relatively cooperative labor relations in Germany. For example, unions and works councils are represented on boards of large companies and thereby involved in managerial decisions.

The approach we develop in this paper differences out the unobserved skills in workers' chosen occupations. This helps solve the econometric selection problem without recurring to parametric assumptions on unobservables. It does, however, come at the cost of not identifying the full population distribution of skills. In the counterfactual analyses in this paper, we therefore condition on observed occupation choices. Other papers, by contrast, assume a static full distribution of skills (e.g., Gabaix and Landier, 2008; Hsieh et al., 2019) to study the effects of important changes in the U.S. economy on the allocation of talent and earnings. One promising avenue for further research is to combine these two approaches, and to obtain micro-identified levels and changes of skills across occupations. With efforts underway to link direct survey measures of skills to administrative records, the data requirements will be met in the near future. Our framework will provide a good starting point to model entire careers. The result would certainly promise to answer key economic questions and allow making predictions about future developments such as the further impact of big data and artificial intelligence.

⁴³Baumgarten and Lehwald (2019) provide evidence for the threat of import competition.

⁴⁴We find few switches of occupations to be systematically associated with large losses, even if we fill intermittent spells of non-employment.

1.A Further Details on the Data and Empirical Facts

1.A.1 Dataset Construction

We employ the SIAB Scientific Use File for our analyses. The SIAB is a 2% random sample of administrative German social security records spanning the years 1975 until 2014.⁴⁵ It includes employees covered by social security, marginal part-time employment (since 1999), unemployment insurance benefit recipients, and individuals who are officially registered as job-seeking or who are participating in programs of active labor market policies. It excludes the self-employed, civil servants, individuals performing military service, and those not in the labor force. In total, it is representative for 80% of the German workforce.

Most notably, it contains an individual's full employment history, including a time-consistent occupational classifier (up to 2010), the corresponding wage, year of birth, place of work, and education. The data is exact to the day as employers need to notify the employment agency if the employment relationship changes. Therefore, the data is available in a spell structure making it possible to observe the same individual at various employers within a year. Those spells may even overlap as workers can have multiple employment contracts at a time. We transform the spell structure into a yearly panel by identifying the longest spell within a given year and deleting all the remaining spells. This procedure differs from the previous inequality literature employing the SIAB in the German context. For instance, Dustmann, Ludsteck, et al. (2009) aggregate all the information from various spells within a year, adding up all the earnings from multiple employment spells. Since our focus is on occupations, this is impossible to do as one cannot aggregate multiple categorical occupation information. Fortunately, the number of workers in our main sample with more than one spell in a year is negligible (< 1.3%) and so of minor concern.

1.A.1.1 Sociodemographics

Occupations: The detailed 120 occupations (KLDB1988) of our main analysis can be found in Table 1.4. Some parts of the analysis make use of a grouping of these 120 occupations into four major classes in the spirit of Acemoglu and Autor (2011):

1. Managers-Professionals-Technicians (Mgr-Prof-Tech)
2. Sales-Office (Sales-Office)
3. Production-Operators-Craftsmen (Prod-Op-Crafts)
4. Services-Care (Srvc-Care)

⁴⁵ Access to the data is subject to signing a contract with the Research Data Center of the German Federal Employment Agency. See Ganzer et al. (2017) for an up to date documentation of the data. We carried out all the analyses making use of the templates provided by Gaudecker (2014).

Education: The education variable contained in the SIAB suffers from some inconsistencies and missing values as described in Fitzenberger et al. (2006) because this information is not irrelevant for social security contributions. We use Fitzenberger et al.'s imputation to obtain an education variable with three or five distinct distinct outcomes: low (missing or without any postsecondary education), medium (apprenticeship training or high school diploma), and high (university degree).

Table 1.4: Grouping of occupations

Group	SIAB occupation
Managers	Entrepreneurs, managing directors, divisional managers Management consultants, organisers until chartered accountants, tax advisers Members of Parliament, Ministers, elected officials until association leaders, officials
Professionals	Architects, civil engineers Artistic and assisting occupations (stage, video and audio) until performers, professional sportsmen, auxiliary artistic occupations Chemists, chemical engineers until physicists, physics engineers, mathematicians Data processing specialists Economic and social scientists, statisticians until scientists n.e.c. Electrical engineers Home wardens, social work teachers Journalists until librarians, archivists, museum specialists Mechanical, motor engineers Music teachers, n.e.c. until other teachers Musicians until scenery/sign painters Navigating ships officers until air transport occupations Physicians until Pharmacists Social workers, care workers until religious care helpers Soldiers, border guards, police officers until judicial enforcers Survey engineers until other engineers University teachers, lecturers at higher technical schools and academies until technical, vocational, factory instructors
Technicians	Biological specialists until physical and mathematical specialists Chemical laboratory assistants until photo laboratory assistants

Continued on next page

Table 1.4: Grouping of occupations

Group	SIAB occupation
Sales	Electrical engineering technicians until building technicians
	Foremen, master mechanics
	Measurement technicians until remaining manufacturing technicians
	Mechanical engineering technicians
	Other technicians
	Technical draughtspersons
	Bank specialists until building society specialists
	Commercial agents, travellers until mobile traders
	Forwarding business dealers
	Health insurance specialists (not social security) until life, property insurance specialists
	Publishing house dealers, booksellers until service-station attendants
	Salespersons
	Tourism specialists until cash collectors, cashiers, ticket sellers, inspectors
Office	Wholesale and retail trade buyers, buyers
	Cost accountants, valuers until accountants
	Office auxiliary workers
	Office specialists
Production	Stenographers, shorthand-typists, typists until data typists
	Building labourer, general until other building labourers, building assistants, n.e.c.
	Ceramics workers until glass processors, glass finishers
	Chemical laboratory workers until vulcanisers
	Chemical plant operatives
	Drillers until borers
	Electrical appliance fitters
	Electrical appliance, electrical parts assemblers
	Engine fitters
	Farmers until animal keepers and related occupations
	Generator machinists until construction machine attendants
	Goods examiners, sorters, n.e.c.
	Goods painters, lacquerers until ceramics/glass painters
	Iron, metal producers, melters until semi-finished product fettlers and other mould casting occupations
	Locksmiths, not specified until sheet metal, plastics fitters
Machine attendants, machinists' helpers until machine setters (no further specification)	

Continued on next page

Table 1.4: Grouping of occupations

Group	SIAB occupation
	Metal grinders until other metal-cutting occupations
	Metal polishers until metal bonders and other metal connectors
	Metal workers (no further specification)
	Miners until shaped brick/concrete block makers
	Other assemblers
	Packagers, goods receivers, despatchers
	Paper, cellulose makers until other paper products makers
	Paviors until road makers
	Plant fitters, maintenance fitters until steel structure fitters, metal shipbuilders
	Plastics processors
	Sheet metal pressers, drawers, stampers until other metal moulders (non-cutting deformation)
	Sheet metal workers
	Special printers, screeners until printer's assistants
	Spinners, fibre preparers until skin processing operatives
	Steel smiths until pipe, tubing fitters
	Tracklayers until other civil engineering workers
	Turners
	Type setters, compositors until printers (flat, gravure)
	Welders, oxy-acetylene cutters
	Wood preparers until basket and wicker products makers
Operators	Assistants (no further specification)
	Motor vehicle drivers
	Post masters until telephonists
	Railway engine drivers until street attendants
	Stowers, furniture packers until stores/transport workers
	Transportation equipment drivers
	Warehouse managers, warehousemen
Craftsmen	Agricultural machinery repairers until precision mechanics
	Bakery goods makers until confectioners (pastry)
	Bricklayers until concrete workers
	Butchers until fish processing operatives
	Carpenters
	Carpenters until scaffolders
	Cutters until textile finishers
	Dental technicians until doll makers, model makers, taxidermists
	Electrical fitters, mechanics

Continued on next page

Table 1.4: Grouping of occupations

Group	SIAB occupation
	Gardeners, garden workers until forest workers, forest cultivators
	Motor vehicle repairers
	Other mechanics until watch-, clockmakers
	Painters, lacquerers (construction)
	Plumbers
	Roofers
	Room equippers until other wood and sports equipment makers
	Stucco workers, plasterers, rough casters until insulators, proofers
	Telecommunications mechanics, craftsmen until radio, sound equipment mechanics
	Tile setters until screed, terrazzo layers
	Toolmakers until precious metal smiths
	Wine coopers until sugar, sweets, ice-cream makers
Service	Cashiers
	Cooks until ready-to-serve meals, fruit, vegetable preservers, preparers
	Doormen, caretakers until domestic and non-domestic servants
	Factory guards, detectives until watchmen, custodians
	Hairdressers until other body care occupations
	Household cleaners until glass, buildings cleaners
	Housekeeping managers until employees by household cheque procedure
	Laundry workers, pressers until textile cleaners, dyers and dry cleaners
	Others attending on guests
	Restaurant, inn, bar keepers, hotel proprietors, catering trade dealers until waiters, stewards
	Street cleaners, refuse disposers until machinery, container cleaners and related occupations
Care	Dietary assistants, pharmaceutical assistants until medical laboratory assistants
	Medical receptionists
	Non-medical practitioners until masseurs, physiotherapists and related occupations
	Nursery teachers, child nurses
	Nurses, midwives
	Nursing assistants

1.A.1.2 Wages and Wage Growth

Despite being accurately measured as the employer can be punished for incorrect reporting, the contained wage variable has two major drawbacks for our analysis.

Wage imputations: First, wages are top coded, amounting to roughly 12% censored observations for men and 2.4% censored observations for women on average across years in our main sample. We impute top coded wages using a series of tobit imputations as in Dustmann, Ludsteck, et al. (2009) or Card et al. (2013), fitted separately for each year, gender and East-West region. We predict the upper tail of the wage distribution employing controls for five age groups and five education groups as well as their interaction and allow the error variance to vary between age and education groups. Further, we include controls for age (within age groups), a part-time dummy, the mean wage in other years, the fraction of censored wages in other years as well as a dummy if the person was only observed once in his life as in Card et al. (2013).⁴⁶ We use the predicted values $X'\hat{\beta}$ from the Tobit regressions together with the estimated standard deviation $\hat{\sigma}$ to impute the censored wages y^c as follows: $y^c = X'\hat{\beta} + \hat{\sigma}\Phi^{-1}[k + u(1-k)]$, where $u \sim U[0, 1]$, $k = \Phi[(c - X'\hat{\beta})/\hat{\sigma}]$ and c is the main censoring limit.⁴⁷ Analog to Jäger and Heining (2019) we scale up daily to yearly wages by multiplying with 365. We inflate wages to prices as of 2010 and finally smooth wages for every individual using three year moving averages.⁴⁸

Wage break 1983/1984: The second major concern with the wage variable is that the definition of a wage – relevant for social security payments – changed between 1983 and 1984. Prior to 1984, wages did not contain bonuses and one time payments. Afterwards these variable parts of the wage were included. If one does not correct this break, it leads to a spurious increase in inequality between those years. We deal with this break by correcting wages prior to 1984 upwards following Fitzenberger (1999) and Dustmann, Ludsteck, et al. (2009). Their idea is that a worker's rank in the wage distribution between 1984 and 1983 should not have changed much. Additionally, they control for the fact that different percentiles of the wage distribution should be differently affected by the break since workers from higher percentiles are likely to receive higher bonuses. Therefore, they estimate locally weighted regressions, separately for men and women, of an individual's wage ratio in 1983/1984 and 1983/1982 on the rank of a person in the wage distribution. Afterwards, they calculate a correction factor as the difference between the predicted, smoothed values from the two wage ratio regressions and multiply wages prior to the break with that factor.

After this, some wages are corrected above the censoring limit. Dustmann, Ludsteck, et al. (2009) reset these wages back to the censoring limit and impute them in the same way they imputed wages which were above the limit anyway. This, however, is very problematic when analyzing wages within high skill occupations. For instance, by employing this procedure, the amount of censored wages within the Mgr-Prof-Tech group aged [45, 54] increases up to 66% in 1975. Instead of following that approach, we do not reset wages back to the censoring limit if they were corrected above the limit

⁴⁶If this is the case, the mean wage in other years and the fraction of censored wages in other years is replaced by the sample mean.

⁴⁷Accessible at http://fdz.iab.de/en/FDZ_Overview_of_Data/working_tools.aspx.

⁴⁸Not smoothing wages does not change the results but leads to spikes in few small occupations.

but leave them at their break corrected values. We create individual log wage growth as the log wage in year t minus the log wage in year $t - 1$ and set it to missing if the worker was not observed in $t - 1$.

1.A.1.3 Sample Selection

The main dataset is restricted to full time working 25 to 54 year old men. We exclude younger workers so that the vast majority of our sample will have concluded their formal education by the time they enter our sample. We stop relatively early because early retirement programs were very common in Germany, particularly in the 1980s and 1990s. Additionally, we drop workers who left the sample for more than 10 years into non-participation, self employment, or the public sector. Workers without information on the occupation are dropped from the analysis. Additionally, the years 2011 - 2014 are left out as the employment agency's official occupational classification changed in 2011. A crosswalk exists in the data but is not 1:1 so that a clear break in employment and wages by occupation is observable between 2010 and 2011. Furthermore, we drop all spells of workers who ever worked in East Germany as well as permanently foreign workers.⁴⁹ The main sample consists of 5,792,111 worker \times year combinations with 428,326 unique workers. Dropping observations with missing information in $t - 1$ results in 5,217,232 observations. The median worker born in the cohort 1950–1956 (the cohort we potentially observe from 25 to 54) is observed for 24 years.

1.A.1.4 Sample with Imputed Non-Employment Spells

One of our key robustness checks (Section 1.F.1) concerns the role of unemployment and out of labor force spells. For this, we relax the exogeneity assumption for unemployment and out of labor force by imputing the occupation where the worker “would have worked in had he not become unemployed or left the labor force”.⁵⁰ We do the imputation by comparing the (real) wage after a non-employment spell with the wage before the non-employment spell. We then impute the wage while in non-employment as the lower of those two wages adjusted for inflation and set the occupation within this time to the occupation that corresponds to that lower wage.⁵¹ The rationale for this procedure is based on the idea that a worker could always choose the lower paying job but eventually decides to quit employment if he prefers becoming unemployed. Imputing unemployed and out of labor force spells results in 6,170,729 observations. Dropping observations with missing information in $t - 1$ leaves us with 5,710,542 observations.

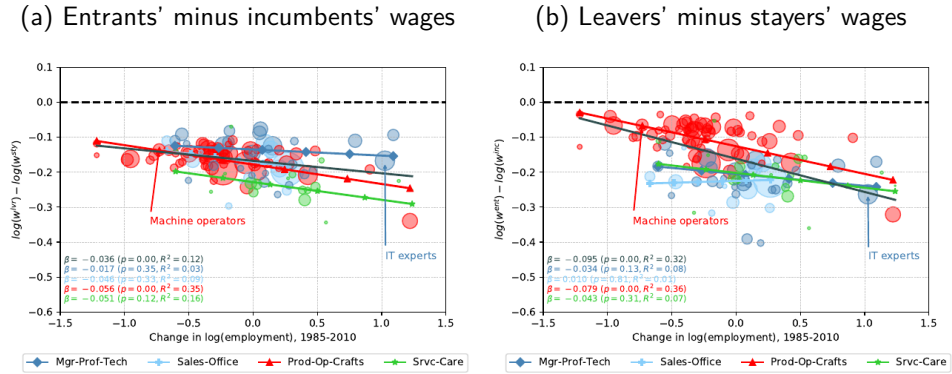
1.A.2 Additional Stylized Facts

⁴⁹That is, workers who are German at some point but foreign at another, are not dropped from the sample. In robustness checks we include the dropped East Germans and foreigners.

⁵⁰Between 1996 and 1998, many workers in occupation 102 “Physicians until Pharmacists” exit the sample and return afterwards as mentioned by Ganzer et al. (2017). We impute those likely erroneously missing observations by setting the occupation to 102 if a worker was in 102 in 1995 and returned in 1999 or 2000 and linearly interpolate the missing wage using the observations in 1995 and 1999/2000. We also drop workers in that group with very low wages between 1988 and 2004 (“Arzt im Praktikum”).

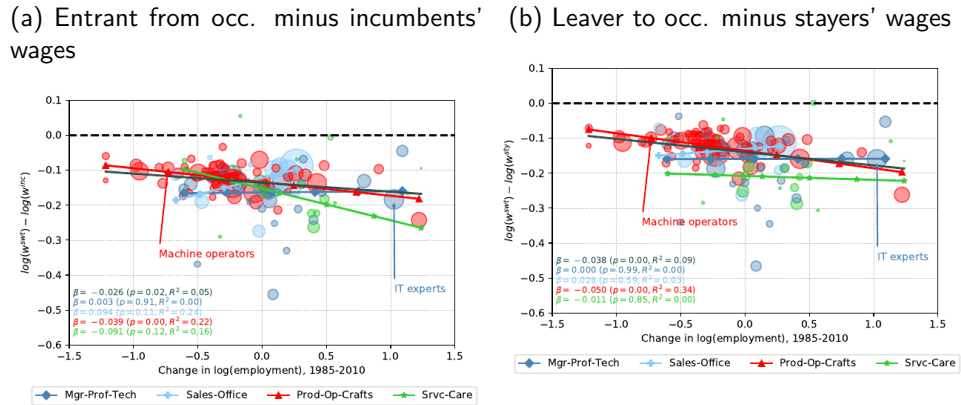
⁵¹As we only fill up spells *between* two employment spells, we therefore treat both unemployment and permanently leaving the labor force without returning to employment as exogenous actions.

Figure 1.10: Selection into and out of occupations, controlling for age and education



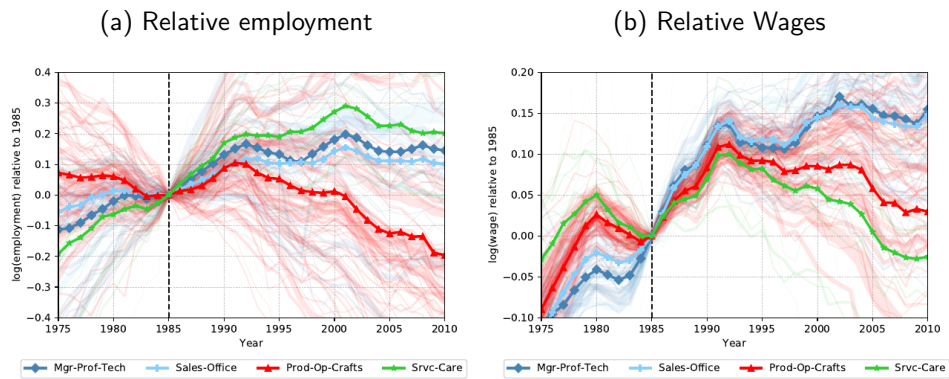
Notes: The vertical axis in Panel 1.10a shows the residual wage of an entrant to an occupation relative to the average wage of incumbents after a regression on age and education dummies. The average is taken across years 1985 until 2010. The vertical axis in Panel 1.10b shows the residual wage of a worker leaving an occupation next period relative to the average wage of stayers after a regression on age and education dummies. The average is taken across years 1985 until 2009 to avoid all workers being leavers at the sample end. The horizontal axis in both panels shows the change of the log number of employed workers within an occupation between 1985 and 2010. One bubble represents one of the 120 detailed occupations in the SIAB SUF. The four groups show an aggregation of these detailed occupations as described in Appendix Table 1.4. Bubble size corresponds to the number of employed workers in an occupation averaged across years 1985 until 2010. Regression lines across all occupations (black) and within the four broad groups (colored) are weighted by the number of employed workers.

Figure 1.11: Selection into and out of occupations, occupational switchers only



Notes: The vertical axis in Panel 1.11a shows the average wage of an entrant to an occupation from another occupation relative to the average wage of incumbents. The average is taken across years 1985 until 2010. The vertical axis in Panel 1.11b shows the average wage of a worker leaving an occupation to another occupation next period relative to the average wage of stayers. The average is taken across years 1985 until 2009 to avoid all workers being leavers at the sample end. The horizontal axis in both panels shows the change of the log number of employed workers within an occupation between 1985 and 2010. One bubble represents one of the 120 detailed occupations in the SIAB SUF. The four groups show an aggregation of these detailed occupations as described in Appendix Table 1.4. Bubble size corresponds to the number of employed workers in an occupation averaged across years 1985 until 2010. Regression lines across all occupations (black) and within the four broad groups (colored) are weighted by the number of employed workers.

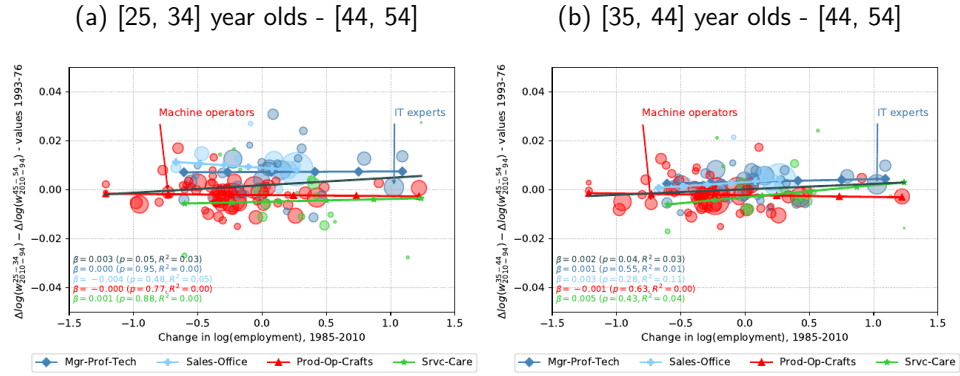
Figure 1.12: Changes in employment and average wages, 1975-2010



Notes: Panel 1.12a shows the log number of workers employed in occupations over time. Panel 1.12b shows the log wage of workers employed in occupations over time. Shaded lines in the background represent the 120 detailed occupations in the SIAB SUF. The four groups show an aggregation of these detailed occupations as described in Appendix Table 1.4. The thickness of a shaded background line corresponds to the number of employed workers in an occupation averaged across years 1985 until 2010.

1.A.3 Constancy of Skill Accumulation

Figure 1.13: Wage growth between age groups



Notes: The figures show a triple difference-in-difference result: how much has wage growth of young (Figure 1.13a) and middle-aged (Figure 1.13b) workers relative to the wage growth of old workers changed after 1993 relative to pre 1993? One bubble represents one of the 120 detailed occupations in the SIAB SUF. The four groups show an aggregation of these detailed occupations as described in Appendix Table 1.4. Bubble size corresponds to the number of employed workers in an occupation averaged across years 1985 until 2010. Regression lines across all occupations (black) and within the four broad groups (colored) are weighted by the number of employed workers.

The figure shows a triple difference-in-difference result: how much has wage growth of young and middle-aged workers relative to the wage growth of old workers changed after 1993 relative to pre 1993? Ideally, the y coordinates of all points would have been close to zero. Despite this not being exactly the case, still 80% of the occupations have absolute differences in growth rates below one percentage point for the 25-34 vs 45-54 comparison; the same holds for 95% of the sample in the 35-44 vs. 45-54 comparison. Importantly, we cannot detect any systematic pattern and there is no clear relation with employment growth, neither overall nor within the four occupation groups. There is one prominent outlier with a very large positive difference. These are medical doctors, who had very high growth rates between 1998 and 2004, a period when low-paid residencies were mandatory for their approbation (“Arzt im Praktikum”).

1.B Theory

1.B.1 Proofs and Derivations

1.B.1.1 Derivation of Equation (1.3)

The goal is to come up with an expression involving the quantity that will eventually be observed— $\Delta w_{i,t}$ —on the left hand side and a sum of integrals on the right hand side. The point at which the notation in (1.3) in the main text does not tell the whole story is that each of these integrals needs to hold constant all wages that are not integrated over; we make this explicit here. To do so, first restate (1.2), explicitly indicating that

$I_{k,i,t}$ is a function of all these potential wages:

$$dw_{i,t} = \sum_{k=1}^K I_k(w_{1,i,t}, \dots, w_{k,i,t}, \dots, w_{K,i,t}) dw_{k,i,t} \quad (1.15)$$

There are $K!$ different permutations of integrating over (1.15); all will lead to the same final result (1.5). We pick one where we start with the potential wage change in occupation 1, holding all other potential wages fixed at their values in period $t-1$:

$$\begin{aligned} & \max(w_{1,i,t}, w_{2,i,t-1}, \dots, w_{K,i,t-1}) - \max(w_{1,i,t-1}, w_{2,i,t-1}, \dots, w_{K,i,t-1}) \\ &= \int_{w_{1,i,t-1}}^{w_{1,i,t}} I_1(w_{1,i,\tau}, w_{2,i,t-1}, \dots, w_{K,i,t-1}) dw_{1,i,\tau} \end{aligned} \quad (1.16)$$

Now continue with occupation 2, holding occupation 1's wage fixed at its value in t and the other occupations' wages at their values in $t-1$:

$$\begin{aligned} & \max(w_{1,i,t}, w_{2,i,t}, w_{3,i,t-1}, \dots, w_{K,i,t-1}) - \max(w_{1,i,t}, w_{2,i,t-1}, w_{3,i,t-1}, \dots, w_{K,i,t-1}) \\ &= \int_{w_{2,i,t-1}}^{w_{2,i,t}} I_2(w_{1,i,t}, w_{2,i,\tau}, w_{3,i,t-1}, \dots, w_{K,i,t-1}) dw_{2,i,\tau} \end{aligned} \quad (1.17)$$

Continue like this so that the conditioning set for occupation k' are the potential wages in t for $k \in \{1, \dots, k'-1\}$ and the potential wages in $t-1$ for $k \in \{k'+1, \dots, K\}$:

$$\begin{aligned} & \max(w_{1,i,t}, \dots, w_{k',i,t}, w_{k',i,t}, w_{k'+1,i,t-1}, \dots, w_{K,i,t-1}) \\ & \quad - \max(w_{1,i,t}, \dots, w_{k'-1,i,t}, w_{k',i,t-1}, w_{k'+1,i,t-1}, \dots, w_{K,i,t-1}) \\ &= \int_{w_{k',i,t-1}}^{w_{k',i,t}} I_{k'}(w_{1,i,t}, \dots, w_{k'-1,i,t}, w_{k',i,\tau}, w_{k'+1,i,t-1}, \dots, w_{K,i,t-1}) dw_{k',i,\tau}. \end{aligned} \quad (1.18)$$

Consider the left-hand side of these terms when summing the K integrals over all occupations. All terms except for $-\max(w_{1,i,t-1}, \dots, w_{K-1,i,t-1}, w_{K,i,t-1})$ from Equation (1.16) and $\max(w_{1,i,t}, w_{2,i,t}, \dots, w_{K,i,t})$, which stems from (1.18) for $k' = K$, drop out. To see this, note that the first term of (1.16) and second term on the left-hand side of (1.17) are the same with opposite signs; this will be the case for any pairs k' and $k'+1$. The two remaining terms are just the observed wage difference $\Delta w_{i,t}$. We thus have:

$$\Delta w_{i,t} = \sum_{k=1}^K \int_{w_{k,i,t-1}}^{w_{k,i,t}} I_k(w_{1,i,t}, \dots, w_{k-1,i,t}, w_{k,i,\tau}, w_{k+1,i,t-1}, \dots, w_{K,i,t-1}) dw_{k,i,\tau} \quad (1.19)$$

The notation of Equation (1.3) in the main text is therefore somewhat imprecise, as each integral with respect to $w_{k,i,\tau}$ in fact holds constant all other wages at their values in either $t-1$ or t . Which of these two values it is depends on the precise order of the integration—(1.19) is just one of $K!$ possibilities—but this is immaterial for the approximation (1.5), which is the point of departure for our empirical analysis.

1.B.1.2 Derivation of Equation (1.5)

Consider the right-hand side of (1.18) and replace the integrand with the linear interpolation (1.4):

$$\begin{aligned}
& \int_{w_{k',i,t-1}}^{w_{k',i,t}} I_{k'}(w_{1,i,t}, \dots, w_{k'-1,i,t}, w_{k',i,\tau}, w_{k'+1,i,t-1}, \dots, w_{K,i,t-1}) dw_{k',i,\tau} \\
& \approx \int_{w_{k',i,t-1}}^{w_{k',i,t}} \left[I_{k',i,t-1} + \frac{I_{k',i,t} - I_{k',i,t-1}}{w_{k',i,t} - w_{k',i,t-1}} (w_{k',i,\tau} - w_{k',i,t-1}) \right] dw_{k',i,\tau} \\
& = I_{k',i,t-1} \Delta w_{k',i,t} + \frac{I_{k',i,t} - I_{k',i,t-1}}{w_{k',i,t} - w_{k',i,t-1}} \left[\frac{1}{2} w_{k',i,\tau}^2 - w_{k',i,t-1} w_{k',i,\tau} \right]_{w_{k',i,t-1}}^{w_{k',i,t}} \\
& = I_{k',i,t-1} \Delta w_{k',i,t} + \frac{1}{2} (I_{k',i,t} - I_{k',i,t-1}) (w_{k',i,t} - w_{k',i,t-1}) \\
& = \bar{I}_{k',i,t} \Delta w_{k',i,t},
\end{aligned} \tag{1.20}$$

where $\bar{I}_{k',i,t} \equiv \frac{I_{k',i,t} + I_{k',i,t-1}}{2}$ is the worker's "average" choice of occupation k' in the two periods. By the derivation in Section 1.B.1.1, summing up over all k gives Equation (1.5):

$$\Delta w_{i,t} = \sum_{k=1}^K \bar{I}_{k,i,t} \Delta w_{k,i,t}$$

1.B.2 Multidimensional as Opposed to One-Dimensional Skills

The easiest way to reject the one-dimensional skill model is to note that the sector-specific wage distributions overlap, e.g., that there exist Mgr-Prof-Tech workers who earn less than some Srvc-Care workers. In the one-dimensional skill model this is impossible because there is a ranking of skill cutoffs above each of which the worker moves to a higher-ranked occupation.

Admittedly, however, the focus of our paper is not in levels of skills but changes of skills and their prices. Therefore, the question is rather whether workers' skill may be described without loss of generality as changing one-dimensionally over the career. This is explored in the following.

1.B.2.1 Theory: Wage Gains When Switching Between Sectors

The multidimensional Roy model postulates $w_{k,i} = \pi_k + s_{k,i}$ as opposed to a one-dimensional skill model a la Cortes (2016) or Acemoglu and Autor (2011), $w_{k,i} = \pi_k + \beta_k s_i$ with s_i a general skill that is priced differently in different sectors according to β_k . The multidimensional skill *change* model flexibly states $\Delta w_{k,i} = \Delta \pi_k + \Delta s_{k,i}$ while the one-dimensional skill change model has the restriction: $\Delta w_{k,i} = \Delta \pi_k + \beta_k \Delta s_i$. This implies that there may be increasing and concave life-cycle profiles of workers in the one-dimensional skill model, idiosyncrasy among these profiles when different workers obtain different shocks, and systematic heterogeneity in the data depending on whether workers switch into more high-skilled sectors (i.e., sectors with a higher skill return β_k) or not. Nonetheless, even in the most general form of this model there are some quite strong empirical restrictions, which we derive now.

To simplify for the moment we continue with constant skill prices $\Delta\pi_k = 0$ and we denote a *general* one-dimensional skill that we condition on of a worker who originates in sector k by $\mathfrak{s}_{k,t-1}$. Notice that for notational brevity we have dropped the individual index i and that k indexes the chosen occupation at time $t-1$, not the (component of multidimensional) skill in that occupation. Suppose the worker starts in sector k with skill $\mathfrak{s}_{k,t-1}$ and stays there or switches either to sector k' or k'' with $\beta_{k''} > \beta_{k'} > \beta_k$. In order to have non-zero employment in all three sectors, we need that $\pi_k > \pi_{k'} > \pi_{k''}$. Therefore, k'' is the highest skill return and skilled sector.⁵² Define the conditions for the choices:

- $k \rightarrow k$: if $\pi_k + \beta_k \mathfrak{s}_{k,t-1} > \pi_{k'} + \beta_{k'} \mathfrak{s}_{k,t-1} > \pi_{k''} + \beta_{k''} \mathfrak{s}_{k,t-1}$ and $\pi_k + \beta_k \mathfrak{s}_{k \rightarrow k,t} > \pi_{k'} + \beta_{k'} \mathfrak{s}_{k \rightarrow k,t} > \pi_{k''} + \beta_{k''} \mathfrak{s}_{k \rightarrow k,t}$. Wage gain: $\Delta w_{k \rightarrow k} = \beta_k (\mathfrak{s}_{k \rightarrow k,t} - \mathfrak{s}_{k,t-1})$.
- $k \rightarrow k'$: if $\pi_k + \beta_k \mathfrak{s}_{k,t-1} > \pi_{k'} + \beta_{k'} \mathfrak{s}_{k,t-1} > \pi_{k''} + \beta_{k''} \mathfrak{s}_{k,t-1}$ and $\pi_{k'} + \beta_{k'} \mathfrak{s}_{k \rightarrow k',t} > \pi_k + \beta_k \mathfrak{s}_{k \rightarrow k',t}$ plus $\pi_{k'} + \beta_{k'} \mathfrak{s}_{k \rightarrow k',t} > \pi_{k''} + \beta_{k''} \mathfrak{s}_{k \rightarrow k',t}$. Wage gain: $\Delta w_{k \rightarrow k'} = \beta_{k'} \mathfrak{s}_{k \rightarrow k',t} - \beta_k \mathfrak{s}_{k,t-1}$.
- $k \rightarrow k''$: Wage gain: $\Delta w_{k \rightarrow k''} = \beta_{k''} \mathfrak{s}_{k \rightarrow k'',t} - \beta_k \mathfrak{s}_{k,t-1}$.

Since the skill levels have to be $\mathfrak{s}_{k \rightarrow k'',t} > \mathfrak{s}_{k \rightarrow k',t} > \mathfrak{s}_{k \rightarrow k,t}$ for the choices in t to be optimal and $\beta_{k''} > \beta_{k'} > \beta_k$ we have a clear ranking of wage gains that we should observe in the data: $\Delta w_{k \rightarrow k''} > \Delta w_{k \rightarrow k'} > \Delta w_{k \rightarrow k}$. This also implies that in this "general" one-dimensional skill model the exogenous mobility assumption is violated, since skill changes conditional on $\mathfrak{s}_{k,t-1}$ are larger depending on where the worker moves.

For one origin occupation k (Srv-Care, say) this is unrestrictive because we can use the observed gains to rank the sectors, i.e., to infer that $\beta_{k''} > \beta_{k'} > \beta_k$ and $\pi_k > \pi_{k'} > \pi_{k''}$ needs to be the case. Using additional origin occupations, k'' (Mgr-Prof-Tech) this does become a restriction as the one-dimensional skill model for a given skill (wage) $\mathfrak{s}_{k'',t-1}$ now prescribes the same ranking in terms of wage gains by destination:

$$\begin{aligned} \Delta w_{k'' \rightarrow k''} &= \beta_{k''} \mathfrak{s}_{k'' \rightarrow k'',t} - \beta_{k''} \mathfrak{s}_{k'',t-1} > \\ &> \Delta w_{k'' \rightarrow k'} = \beta_{k'} \mathfrak{s}_{k'' \rightarrow k',t} - \beta_{k''} \mathfrak{s}_{k'',t-1} > \\ &> \Delta w_{k'' \rightarrow k} = \beta_k \mathfrak{s}_{k'' \rightarrow k,t} - \beta_{k''} \mathfrak{s}_{k'',t-1}, \end{aligned}$$

since $\mathfrak{s}_{k'' \rightarrow k'',t} > \mathfrak{s}_{k'' \rightarrow k',t} > \mathfrak{s}_{k'' \rightarrow k,t}$ for the choices in t to be optimal and $\beta_{k''} > \beta_{k'} > \beta_k$. Similarly, we expect in the data that $\Delta w_{k' \rightarrow k''} > \Delta w_{k' \rightarrow k'} > \Delta w_{k' \rightarrow k}$. We can also condition on the destination sector. Fixing k with skill $\mathfrak{s}_{k,t}$ we get

$$\begin{aligned} \Delta w_{k \rightarrow k} &= \beta_k \mathfrak{s}_{k,t} - \beta_k \mathfrak{s}_{k \rightarrow k,t-1} > \Delta w_{k' \rightarrow k} = \beta_k \mathfrak{s}_{k,t} - \beta_{k'} \mathfrak{s}_{k' \rightarrow k,t-1} > \\ &> \Delta w_{k'' \rightarrow k} = \beta_k \mathfrak{s}_{k,t} - \beta_{k''} \mathfrak{s}_{k'' \rightarrow k,t-1}, \end{aligned}$$

since $\mathfrak{s}_{k \rightarrow k,t-1} < \mathfrak{s}_{k' \rightarrow k,t-1} < \mathfrak{s}_{k'' \rightarrow k,t-1}$ for the choices in $t-1$ to be optimal and $\beta_{k''} > \beta_{k'} > \beta_k$. Similarly, $\Delta w_{k \rightarrow k'} > \Delta w_{k' \rightarrow k'} > \Delta w_{k'' \rightarrow k'}$ and $\Delta w_{k \rightarrow k''} > \Delta w_{k' \rightarrow k''} > \Delta w_{k'' \rightarrow k''}$.

⁵²These considerations directly follow Cortes (2016). He also shows that there exist unique cutoffs s' and s'' determined by indifference that span mutually exclusive and exhaustive intervals $(-\infty, s']$, $(s', s'']$, and $(s'', \infty]$ of skills within which individuals choose work in k , k' , and k'' , respectively.

We therefore obtain the following empirical restrictions from the one-dimensional (general as opposed to specific) skill model:

1. For any given origin sector k' and skill $s_{k',t-1}$, there is a fixed ranking of wage gains by destination sector. That is, the size ordering of wage gains $\{\Delta w_{k',1}, \dots, \Delta w_{k' \rightarrow k}, \dots, \Delta w_{k' \rightarrow K}\}$ does not depend on k' .
2. For any given destination sector k' and skill $s_{k',t}$, there is a fixed ranking of wage gains by origin sector. That is, the size ordering of wage gains $\{\Delta w_{1,k'}, \dots, \Delta w_{k \rightarrow k'}, \dots, \Delta w_{K,k'}\}$ does not depend on k' and it is exactly the inverse ordering of (1.), the ordering of destination sectors, in the running index $1, \dots, K$.

We have abstracted from changes of skill prices in this argument. If sectors' skill intensities and wage ranks do not invert, i.e., the qualitative ranking $\beta_{k''} > \beta_{k'} > \beta_k$ and $\pi_k > \pi_{k'} > \pi_{k''}$ remains stable over time, which is strongly supported in the data, evolving skill prices do not affect the above results. The reason is again by revealed preference: conditioning on the same origin sector and skills, in order for his decision to be optimal, a worker switching into k'' has to have higher skill gains than a worker switching into k' . Since the worker switching into k'' could always switch into k' and have higher wage gains than the worker who actually decides to switch into k' , the former worker's realized wage gain must be higher than the latter. If sectors' skill intensities or wage ranks had inverted in the data, we could always condition on sub-periods of our sample where they did not do that. Therefore, the empirical restrictions (1.) and (2.) from the one-dimensional skill model persist for the case of generally evolving skill prices over time.

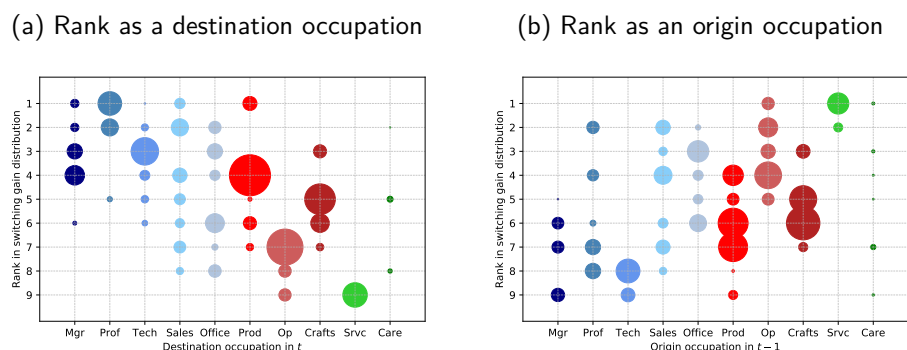
A couple more features to notice:

- These restrictions do not depend on whether the skill changes arise from systematic accumulation or idiosyncratic shocks. The distribution of shocks does also not have to be known and can differ conditional on origin or destination sector. The key assumption that generates the restrictions is that the skill accumulation or shocks are general (one-dimensional), not sector specific!
- This model does not restrict that workers can learn more in some sectors than in others, e.g., that skill growth in Mgr-Prof-Tech is *on average* larger than in Srvc-Care. It is just that all skill growth is general (one-dimensional), not specific.
- One very helpful feature here is also that we can condition on the wage in the origin (1.) or destination (2.) sector and thus perfectly fix the worker's initial $s_{k',t-1}$ or final skill, since skills are indeed the same in each sector!

1.B.2.2 Evidence: Gains from Switching Into or Out of Ten Broad Occupations

Figure 1.14a shows the rank of unconditional wage gains by each out of ten broad destination occupations. Restriction (1.) of the one-dimensional skill model predicts that there is a consistent ranking of these wage gains regardless of the origin occupations. We have already pre-ordered them using some prior knowledge and we see some of this in the figure, whereby wage gain ranks decline from the top-left of the graph to the bottom-right. In particular, workers moving into Mgr or Prof occupations tend to have highly ranked wage gains whereas workers moving into Srvc or Care occupations have among the lowest wage gains.

Figure 1.14: Gains from switching



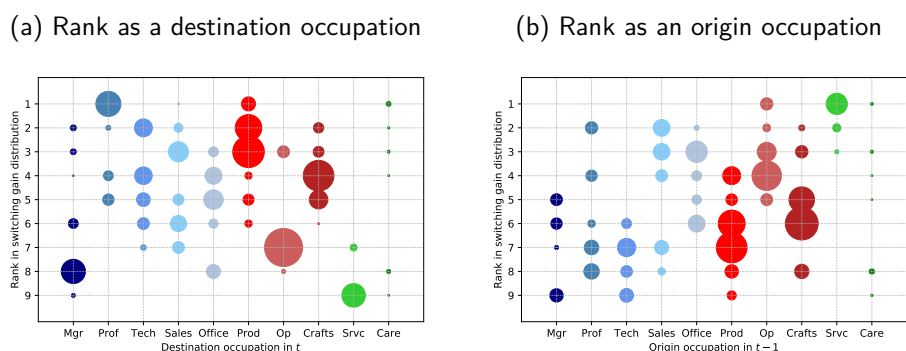
Notes: The ten groups are based on an aggregation of detailed occupations in the SIAB SUF as described in Appendix Table 1.4. Bubbles show the rank (1 is highest, 9 is lowest) of an occupation in the distribution of average wage gains from switching. Bubble size corresponds to the amount of switchers. Panel 1.14a shows the rank in (unconditional) wage gains from all occupations when the one on the x-axis is the destination occupation. Panel 1.14b shows the rank in (unconditional) wage gains into all occupations when the one on the x-axis is the origin occupation.

However, there is also a substantial amount of heterogeneity in these wage gain ranks. For example, it really depends on where a worker starts out from whether he has highly ranked wage gains moving into Sales or Office occupations, ranging from top gains (bubble at 1) to almost bottom (bubble at 8). Specifically, the highest gains of switching into Sales are for Prof and Office (bubbles at 1 and 2), while three out of the four lowest ranked gains are for the Prod, Op, or Crafts occupations. This makes sense if we think that workers acquire *relatively* little Sales-relevant skills (e.g., communication and persuasion) when working in Prod-Op-Crafts occupations, or if workers who choose Prod-Op-Crafts were initially endowed with relatively little Sales skills. In contrast, Tech is a high-gain destination for Prod-Op-Crafts workers (bubbles at 2 and 3) whereas Sales and Office (bubbles at 4 and 5) have lower gains moving into Tech. These different rankings of gains can also broadly be seen in our estimated skill accumulation Table 1.19 below, and they reflect the fact that workers in different occupations have different specific skills.

Figure 1.14b shows the inverse graph to what we just discussed, that is, the wage gains by origin occupations. Restriction (2.) of the one-dimensional skill model predicts that those are also consistently ranked and that the ranking is the inverse of the wage gains by destination. Indeed, we see that the gains in Figure 1.14b tend to rise from the bottom left to top right inversely to Figure 1.14a, and they tend to be lowest for workers moving out of Mgr-Prof-Tech occupations and highest for switchers out of Srv and Care. However, we would expect the gains/losses from switching to be consistently ranked, i.e., movers out of Mgr having always the highest losses for all destination occupations, professionals always having the second-highest losses, up until workers leaving Srv having the highest gains no matter what is the destination occupation. This is clearly not the case in Figure 1.14b as there is again a substantial amount of heterogeneity in ranks depending on the destination and in fact in all of the ten

occupations other than Srvc.⁵³ The gains are also far from perfectly inverted.

Figure 1.15: Gains from switching (residual)



Notes: The ten groups are based on an aggregation of detailed occupations in the SIAB SUF as described in Appendix Table 1.4. Bubbles show the rank (1 is highest, 9 is lowest) of an occupation in the distribution of average wage gains from switching. Bubble size corresponds to the amount of switchers. Panel 1.15a shows the rank in residual wage gains from all occupations when the one on the x-axis is the destination occupation. The residual wage growth holds origin wages constant, i.e., the residual from a regression of $\Delta w_{i,t}$ between t and $t-1$ on a worker's previous wage in $t-1$. Panel 1.15b shows the rank in residual wage gains into all occupations when the one on the x-axis is the origin occupation. The residual wage growth holds destination wages constant, i.e., the residual from a regression of $\Delta w_{i,t}$ between t and $t-1$ on a worker's current wage in t .

The results we just discussed are qualitatively the same when controlling for wage in the origin and destination occupations, respectively in Figure 1.15's Panels a and b, as prescribed by the theory in the previous section. The dispersion of gains is in fact more heterogeneous. That is, conditioning on origin skill brings out even more that there are differing rankings of destination sectors by origin occupation. Conditioning on destination skill brings out even more that there are differing rankings of origin sectors by destination occupation. Both of these results point to skills that are specific for different origin-destination combinations.

When controlling for age and other observables (not depicted) the results are again the same. We therefore conclude from this evidence that the one-dimensional skill (change) model can be a reasonable approximation of wage gains in some circumstances but overall it is rejected in the data with its heterogeneity of wage gain ranks. This matters for our results in the paper because it generates the evidence for direct occupation switchers in Table 1.1, and more broadly the fact in Figure 1.3 that both entrants and leavers in any occupation earn less than the respective incumbents and

⁵³For example, the gains of Prof moving into Mgr or Tech (bubbles in Prof column at 7 and 8) are low while the gains of Prod (bubbles at 4 and 5) or Crafts (bubbles at 3 and 5) moving into those occupations are higher. This is consistent with the one-dimensional skill model where Prof are highly-ranked occupations, moving from which can hardly constitute an improvement, whereas Prod and Crafts are rather middle-ranked occupations. However, the gains of switching from Prof as an origin are rather dispersed; especially moving into Office (bubble at 2) or Sales (bubble at 4) are quite high. In contrast, Prod (bubbles at 5 and 9) or Crafts (bubbles at 6 and 7) as an origin occupation have among the lowest gains moving into Office and Sales. Again, this is consistent with Prof workers having *relatively* high skills in Sales and Office compared to, for example, Prod and Crafts workers.

stayers.

1.B.3 Extensions of the Model

1.B.3.1 Costs of Switching Occupations

Switching occupations, which is often requires moving to a different employer and / or city, may be associated with pecuniary and non-pecuniary non-wage costs, e.g., financial expenses and psychological stress of moving house to be close to the new job, which are potentially as important as wage costs (Dix-Carneiro, 2014; Artuç and McLaren, 2015; Cortes and Gallipoli, 2017).

For simplicity, assume that there is a fixed cost of switching occupations $c > 0$ which is incurred if and only if a worker is moving to a different occupation. Define the modified variable entering the decisions of the worker as $w_{k',i,t}^*$ as:

$$w_{k',i,t}^* = \begin{cases} w_{k',i,t} & \text{if } I_{k',i,t-1} = 1 \\ w_{k',i,t} - c & \text{if } I_{k',i,t-1} = 0 \end{cases} \quad (1.21)$$

That is, a worker switches occupations only if the resulting wage is at least c higher than in his origin occupation. Equipped with the decision-relevant wage $w_{k',i,t}^*$, we can make the derivations corresponding to Section 1.B.1.1. In particular, workers are indifferent in terms of $w_{k',i,t}^*$ at the switch point but realized wages will now discretely jump at that point by a function of c .

The equivalent of (1.19) is:

$$\Delta w_{i,t}^* = \sum_{k=1}^K \int_{w_{k,i,t-1}^*}^{w_{k,i,t}^*} I_k(w_{1,i,t-1}^*, \dots, w_{k-1,i,t-1}^*, w_{k,i,\tau}^*, w_{k+1,i,t}^*, \dots, w_{K,i,t}^*) dw_{k,i,\tau}^* \quad (1.22)$$

Linearly approximating the integral, we end up with $\Delta w_{i,t}^* = \sum_{k=1}^K \bar{I}_{k,i,t} \Delta w_{k,i,t}^*$. This can be rearranged in terms of realized wages similar to Section 1.B.3.2 below:

$$\Delta w_{i,t} = \sum_k \bar{I}_{k,i,t} \Delta w_{k,i,t} + \sum_k \bar{c}_{k,i,t} \Delta I_{k,i,t} \quad (1.23)$$

where $\bar{c}_{k,i,t}$ is the notation for i 's average costs incurred in the two periods t and $t-1$ because of switching into or out of occupation k . The observed wage difference $\Delta w_{i,t}$ for switchers (i.e., when $\Delta I_{k,i,t} \neq 0$) will be larger than the difference for the decision variable $\Delta w_{i,t}^*$.

We deliberately employed the notation that is typically used in measurement error models in econometrics: We would love to observe $w_{i,t}^*$, all we see in the data is $w_{i,t}$. While there maybe some hope for estimating a basic version with strong restrictions on c (like it being a constant across all occupations or a fraction of the previous wage), we view any restrictions that may have empirical bite as too strong. They will likely vastly differ across regions, urban/rural areas, the "distance" between occupations, whether the same employer offers different occupations, and individual characteristics like family composition. There is no chance to observe a meaningful subset of such factors in our data. We hence treat c as unobservable and it will enter the error term of the estimation Equation (1.9) much in the same way that $u_{k,i,t}$ does. Of course, it will not have mean zero anymore. In particular, it will exacerbate the correlation between the unobservables and the period- t -choice, leading to upward-biased skill accumulation

coefficients of workers switching occupations but otherwise not affect the estimates. We extensively examine switching costs in our Monte Carlo simulations, Section 1.C.4 of this Appendix.

1.B.3.2 Non-Pecuniary Benefits

In the discussion of the main text, individuals are myopic and maximize their current wages. In this section, we show how the model can be reinterpreted to accommodate non-pecuniary valuations of different occupations. Suppose that the utility of worker i in occupation k at time t is:

$$U_{k,i,t} = w_{k,i,t} + V_{k,i,t} \text{ with} \quad (1.24)$$

$$V_{k,i,t} = X'_{i,t-1} \Psi_{k,t} + \varepsilon_{k,i,t}, \quad (1.25)$$

where $V_{k,i,t}$ is occupation k 's amenity value; it may vary across workers. The discussion here largely follows Böhm (2019). See, for example, Lee and Wolpin (2006) for a full-fledged structural model that incorporates both, forward-looking behavior and non-pecuniary amenities.

In particular, the vector $\Psi_{k,t}$ contains an intercept as well as occupation-specific mappings from $X_{i,t-1}$ to utility. That is, the non-pecuniary or continuation value of each occupation k will differ by workers' characteristics like age or education in practice. We further let idiosyncratic occupation valuations $\varepsilon_{k,i,t}$ be mean zero and independent across individuals. They may be correlated across occupations for a given individual. Finally, notice already that only relative values $V_{k,i,t}$, and thus the parameters $\Psi_{k,t}$ compared to a chosen reference occupation, will be identifiable from workers' observed choices and wages.

With Equation (1.24) and utility maximization at hand, we can make the derivations corresponding to the main text. In particular, the marginal relationship (1.2) holds for utilities and we arrive at $\Delta U_{i,t} = \sum_{k=1}^K \int_{U_{k,i,t-1}}^{U_{k,i,t}} I_{k,i,\tau} dU_{k,i,\tau}$. That is, workers are still indifferent in utility terms at the switch point. We cannot, however, separately write $\Delta w_{i,t} = \sum_{k=1}^K \int_{w_{k,i,t-1}}^{w_{k,i,t}} I_{k,i,\tau} dw_{k,i,\tau}$ (and $\Delta V_{i,t} = \sum_{k=1}^K \int_{V_{k,i,t-1}}^{V_{k,i,t}} I_{k,i,\tau} dV_{k,i,\tau}$) anymore, since wages may indeed discretely jump at the indifference point. This is because indifference is now determined by utility and it may be the case that, at the point of switching, wages are strictly higher and amenity/continuation values are strictly lower in the destination than in the origin occupation, or vice versa.

Linearly approximating the integral again, we arrive at an augmented Equation (1.5) for the change of utility between two periods:

$$\Delta U_{i,t} = \Delta w_{i,t} + \Delta V_{i,t} = \sum_{k=1}^K \bar{I}_{k,i,t} \Delta U_{k,i,t} = \sum_{k=1}^K \bar{I}_{k,i,t} (\Delta w_{k,i,t} + \Delta V_{k,i,t}), \quad (1.26)$$

where $\bar{I}_{k,i,t} \equiv 1[\max_{k'=1,\dots,K} \{U_{k',i,t}\} = U_{k,i,t}] = 1[U_{k,i,t} \geq U_{k',i,t} \forall k' \neq k]$ is now the definition of the choice indicator for occupation k .

The intuition here is also parallel to Equation (1.5): if a worker stays in his occupation, his realized utility gain is the change of his potential utility in that occupation. That is, $\Delta U_i = \Delta U_{ki}$ if $I_{k,i,t} = I_{k,i,t-1} = 1$, which is not an approximation. If the worker switches (e.g., occupations k' to k , $I_{k,i,t-1} = 1, I_{k,i,t} = 1$), he obtains part of the origin occupation's utility gain (or loss) as well as part of the destination occupation's utility gain, set to half-half by the approximation (i.e., $\Delta U_{i,t} = \frac{1}{2} \Delta U_{k',i,t} + \frac{1}{2} \Delta U_{k,i,t}$). The arguments of the main text why the approximation error is negligible apply.

We want to solve Equation (1.26) for $\Delta w_{i,t}$, which is observable in the data. Consider $V_{i,t} = \sum_k I_{k,i,t} V_{k,i,t}$ and then write

$$\begin{aligned}\Delta V_{i,t} &= \sum_k I_{k,i,t} V_{k,i,t} - \sum_k I_{k,i,t-1} V_{k,i,t-1} \\ &= \sum_k \bar{I}_{k,i,t} \Delta V_{k,i,t} + \sum_k \bar{V}_{k,i,t} \Delta I_{k,i,t},\end{aligned}$$

with $\bar{V}_{k,i,t} \equiv \frac{1}{2}(V_{k,i,t-1} + V_{k,i,t})$ and $\Delta I_{k,i,t} \equiv I_{k,i,t} - I_{k,i,t-1}$. Inserting this into Equation (1.26), the realized wage growth of individual worker i in the generalized Roy model becomes:

$$\Delta w_{i,t} = \sum_k \bar{I}_{k,i,t} \Delta \pi_{k,t} + \sum_k \bar{I}_{k,i,t} \Delta s_{k,i,t} - \sum_k \bar{V}_{k,i,t} \Delta I_{k,i,t}. \quad (1.27)$$

This result firstly has a purely pecuniary part as in Equation (1.5) of the main text: if a worker stays in his occupation, his wage gain is the potential wage change (i.e., price growth and skill accumulation) in that occupation. If the worker switches, he obtains half of the origin's as well as half of the destination's potential wage change. Similar to (1.5), Equation (1.27) accommodates endogenous switches, which in this case may be due to changes in amenity/continuation values in addition to changes in potential wages.

The third summand on the right of Equation (1.27) is then the intuitive extension of a purely pecuniary/static model: with optimal choices, a worker's observed wage growth is the change in the potential wage of his chosen occupations minus the utility gain (loss) from the behavioral response of switching occupations. That is, if a utility-optimizing worker chooses to switch occupations (e.g., from k' to k so that $\Delta I_{k,i,t} = 1$ and $\Delta I_{k',i,t} = -1$), we observe lower wage growth than the change in relevant potential wages when he gains amenities or net present value of future earnings (i.e., $\bar{V}_{k,i,t} > \bar{V}_{k',i,t}$ and thus $\sum_k \bar{V}_{k,i,t} \Delta I_{k,i,t} > 0$) via the move. Vice versa, we observe higher wage growth than the potential wage changes when he moves to a less desirable occupation in these respects ($\sum_k \bar{V}_{k,i,t} \Delta I_{k,i,t} < 0$).

Notice in Equation (1.27) it is the *average* non-pecuniary value over both periods $\bar{V}_{k,i,t}$ that the worker is moving into which matters for wage changes. For a switcher from k' to k ,

$$\sum_k \bar{V}_{k,i,t} \Delta I_{k,i,t} = \bar{V}_{k,i,t} - \bar{V}_{k',i,t} = \frac{1}{2}(V_{k,i,t} - V_{k',i,t}) + \frac{1}{2}(V_{k,i,t-1} - V_{k',i,t-1}),$$

conditional on wage gains associated with average choices, moving into the currently high-value occupation (i.e., $V_{k,i,t} - V_{k',i,t} > 0$) is offset with lower wage growth. But also moving into a occupation that last period carried high value ($V_{k,i,t-1} - V_{k',i,t-1} > 0$) is associated with lower wage growth because it implies that the worker was compensated last period for working in the low-value occupation, which now falls away with the switch. Both of these factors enter equally into the wage Equation (1.27). Hence one cannot distinguish them empirically and only identify the average value over the two periods. As is obvious from Equation (1.27), one can also not distinguish between the amenity and the continuation value considerations but only estimate a joint parameter $\bar{V}_{k,i,t}$. Finally, notice when the worker makes no switch, the value considerations do not come into play at all (i.e., $\sum_k \bar{V}_{k,i,t} \Delta I_{k,i,t} = 0$) and the changing wage is just the changing skill price plus skill accumulation.

We now discuss empirical implementation for different versions of Equation (1.25). First, if non-pecuniary values are constant such that $\bar{V}_{k,i}$ does not carry a time index, they will be simply incorporated in the skill accumulation parameters. That is, since $\sum_{k'=1}^K \bar{I}_{k,i,t} I_{k',i,t-1} \cdot X'_{i,t-1} \Gamma_{k',k}$ in Equation (1.9) is a fully interacted model of all task choice combinations and worker observables, it absorbs the term $\sum_k \bar{V}_{k,i} \Delta I_{k,i,t}$. Because of this, also in the time-varying case, average $\bar{V}_{k,i,t}$ parameters can only be identified relative to their base period values. Our main estimation specification therefore already controls for general time-invariant non-pecuniary values as well as forward-looking considerations of occupation choice (with the interpretation of the parameter estimates $\hat{\Gamma}_{k',k}$ adjusted accordingly).

If instead non-pecuniary values are time-varying, we first of all note again that only $\bar{V}_{k,i,t}$ relative to a reference occupation can be identified. The mechanical reason is that the $\Delta I_{k,i,t}$ sum to zero over all K , and thus one of them has to be left out of the estimation due to multicollinearity. The economic intuition is that we can use choices and wages to identify relative utilities but not their levels. Other than that, it is straightforward to introduce a full set of task choice changes into estimation Equation (1.9) and also interact them with worker characteristics. That is, $\bar{\Psi}_{k,t}$ in augmented regression

$$\begin{aligned} \Delta w_{i,t} = & \sum_{k=1}^K \bar{I}_{k,i,t} \cdot \Delta \pi_{k,t} + \sum_{k'=1}^K \bar{I}_{k,i,t} I_{k',i,t-1} \cdot X'_{i,t-1} \Gamma_{k',k} \\ & - \sum_{k'=1}^K \Delta I_{k,i,t} \cdot X'_{i,t-1} \bar{\Psi}_{k,t} + v_{i,t}, \end{aligned} \quad (1.28)$$

identifies the average between time t and $t-1$ amenity/continuation values in k relative to a reference occupation and to the base period by age group.

The reason for why (1.28) is identified, even with time-varying valuations, is the above-discussed fact that moving into current as well as past $\bar{V}_{k,i,t}$ both matter equally for wage changes. Therefore, this contribution to wage growth is only via the changing sorting $\Delta I_{k,i,t}$ into average non-pecuniary values, whereas the contribution to wage growth from changing skill prices and skill accumulation is only via the average sorting $\bar{I}_{k,i,t}$. The estimation method is thus robust to both changing potential wages and non-pecuniary values over time.

Finally, conditional on specification (1.28), an additional (average) error term $\bar{\varepsilon}_{k,i,t}$ in (1.25) does not much affect the estimates, which to some extent parallels the limited confounding role of idiosyncratic skill shocks in the main text (see Böhm, 2019, for more detailed discussion of the idiosyncratic non-pecuniary error term). We estimate (1.28) in Figure 1.47 and report very similar skill price estimates to our main results. For younger workers, the non-pecuniary values of Mgr-Prof-Tech, Sales-Office, and Srvc-Care have modestly declined compared to Prod-Op-Crafts over the sample period.

1.B.3.3 Learning About Skills

In the rest of the paper we have assumed that, aside from skill prices, all changes in individuals' wages over time are due to systematic skill accumulation and idiosyncratic skill shocks. In this section, we show that the model's interpretation can be widened to include imperfect information about skills and employer learning over time in addition

to skill accumulation.⁵⁴

Suppose that, as in the employer learning literature (e.g., Altonji and Pierret, 2001; Gibbons, Katz, et al., 2005), information about skills is imperfect. Each period an additional noisy signal of the worker's productivity arrives; employers form expectations about skills based on this as well as on all past observable information. Expectations are rational in the sense that employers' beliefs are correct on average. Information is symmetric, employers are competitive, all market participants are risk neutral, and a spot market for labor exists.

In this setup, workers' potential log wages in each occupation equal their expected productivity conditional on all available information:

$$w_{k,i,t} = \pi_{k,t} + E_t(s_{k,i,t}) \quad \forall k \in \{1, \dots, K\}, \quad (1.29)$$

where E_t indicates that we are conditioning on all the information available in t . We assume that workers maximize their log incomes by choosing the occupation in which they earn the highest wage. This yields a modified version of Equation (1.5) for observed wage growth over time:

$$\Delta w_{i,t} = \sum_{k=1}^K \bar{I}_{k,i,t} \Delta \pi_{k,t} + \sum_{k=1}^K \bar{I}_{k,i,t} \Delta E_t(s_{k,i,t}), \quad (1.30)$$

where $\Delta E_t(s_{k,i,t}) \equiv E_t(s_{k,i,t}) - E_{t-1}(s_{k,i,t-1})$ and the linearity in logs allows us to swap the summation, first differencing, and expectations operators. The analogue to the skill accumulation Equation (1.7) becomes:

$$\Delta E_t(s_{k,i,t}) = \sum_{k'=1}^K I_{k',i,t-1} \cdot X'_{i,t-1} \Gamma_{k',k} + u_{k,i,t}. \quad (1.31)$$

The only differences lie on the left-hand side—where changes in expected skills replace changes in actual skills—and in the interpretation of the innovations $u_{k,i,t}$, which now represents an update of employers' expectations about individual i 's skill. These changes are immaterial for our estimation strategy; our results remain valid under a basic model of employer learning about skills as an alternative or in addition to systematic skill accumulation and idiosyncratic skill shocks.

1.B.4 Occupation-Specific Fixed Effects as an Alternative Approach

In this section, we examine the occupation-specific fixed effects approach for estimating skill prices as an alternative to our method. We show that under a flexible model of skill accumulation, this approach requires controlling for workers' whole history of occupation-specific experience or, more feasibly, extending the fixed effects to being occupation-stint specific. A base period or some other restriction on the skill accumulation are needed. With idiosyncratic skill shocks, an endogeneity bias emerges that is due to the fixed effects themselves. The results from the Monte Carlo simulations in Section 1.C support our analytical arguments.

⁵⁴Groes et al. (2014) do it the other way around; they set up their theoretical model as employer learning but then clarify that it could alternatively be "shocks to workers' ability" p. 5 Groes et al., 2014.

Several papers have used fixed effects approaches in order to address worker heterogeneity when estimating skill prices (e.g., Cortes, 2016; Cavaglia and Etheridge, 2017).⁵⁵ To be specific, consider Cortes' time-varying model for the potential wage of individual i in occupation k at time t :

$$w_{k,i,t} = \pi_{k,t} + s_{k,i,t} = \pi_{k,t} + X'_{i,t}\Gamma_k + \eta_{k,i}. \quad (1.32)$$

The changing characteristics vector $X_{i,t}$ can increase skills differentially with age or experience across occupations according to Γ_k . In addition, $\eta_{k,i}$ are occupation-specific time-invariant skill levels, which will be introduced into the regression by individual-occupation specific fixed effects. Cortes (2016) and Cavaglia and Etheridge (2017) interchangeably call these occupation- or sector-*spell* fixed effects, which is why we instead use the term 'stint' for a worker's self-contained stay (i.e., without switches in between) in a given occupation below. Consistent with (1.32), Cortes' estimation equation (8) in our notation is:⁵⁶

$$w_{i,t} = \sum_{k=1}^K I_{k,i,t}\pi_{k,t} + \sum_{k=1}^K I_{k,i,t}\eta_{k,i} + \sum_{k=1}^K I_{k,i,t} \cdot X'_{i,t}\Gamma_k + \varphi_{i,t}, \quad (1.33)$$

where we have added the idiosyncratic error term $\varphi_{i,t}$. In the following, we examine under what conditions estimation of Equation (1.33) may then identify the correct skill prices.

1.B.4.1 Systematic Skill Accumulation

We start by assuming that, as in Cortes (2016) or Cavaglia and Etheridge (2017), $\varphi_{i,t}$ is simply measurement error and thus not decision-relevant (exogenous mobility assumption of fixed effects approaches discussed in the main text). The skill accumulation becomes

$$\Delta s_{k,i,t} = \sum_{k'=1}^K I_{k',i,t-1} \cdot \Gamma_{k',k}, \quad (1.34)$$

where, compared to Equation (1.7), we omit for now the $X_{i,t-1}$ -specificity of the skill accumulation function as another simplifying assumption and thus $\Gamma_{k',k}$ is a scalar. Writing this out from when the worker joined the labor market at time $t_{i,0}$ gives

$$s_{k,i,t} = \eta_{k,i} + \sum_{k'=1}^K [I_{k',i,t-1} + \dots + I_{k',i,t_{i,0}}] \cdot \Gamma_{k',k} = \eta_{k,i} + \sum_{k'=1}^K \sum_{\tau=t_{i,0}}^{t-1} I_{k',i,\tau} \cdot \Gamma_{k',k}, \quad (1.35)$$

for $t \geq t_{i,0}$ and $\eta_{k,i}$ the initial skill endowments of i in k when he joins the labor market. Therefore, if we are willing to assume that skill accumulation occurs similarly in each occupation of origin ($\Gamma_{k',k} = \Gamma_k, \forall k', k$), this simplifies to $s_{k,i,t} = \eta_{k,i} + (t -$

⁵⁵In more broadly related settings, Combes et al. (2008) estimate city wage premia, taking into account sorting across locations. Analyzing variation over the business cycle, Solon et al. (1994) account for skill selection into the labor market market, while McLaughlin and Bils (2001) examine skill selection across sectors.

⁵⁶Similar to us, Cortes (2016) uses ten year age bins in $X_{i,t-1}$, allowing for the convexity of the life-cycle profile parallel to our Equation (1.7).

$t_{i,0}) \cdot \Gamma_k$ and Estimation (1.33) identifies the correct skill prices, initial endowments, and skill accumulation parameters. In this case, $X_{i,t} = t - t_{i,0}$ represents labor market experience (proxied in Cortes (2016) by age dummies) in the estimation. Notice that this specification assumes that labor market experience is not occupation-specific, just that general experience is valued differently in different occupations. In Section 1.B.2 we formally test and reject such a one-dimensional skill model.

The need for a base period or similar restriction: We have argued in the main text that any approach using panel data needs a base period or other fundamental restriction of the skill accumulation function. This is the same in Estimation (1.33) and easiest to see if we simplify it to its essence. First, noting that, because of the individual-occupation-specific fixed effects $\eta_{k,i}$, the changing skill prices $\pi_{k,t}$ are fundamentally identified from wage growth of occupation stayers. Therefore, we can condition on the respective occupation k :

$$w_{k,i,t} = \pi_{k,t} + \eta_{k,i} + X'_{i,t} \Gamma_k + \varphi_{i,t} \quad (1.36)$$

Second, fixed effects estimates are asymptotically equivalent to first differences (and exactly the same in finite samples if $T = 2$). First-differencing gives the effective variation that the occupation-specific fixed effects approach identifies from:⁵⁷

$$\Delta w_{k,i,t} = \Delta \pi_{k,t} + \Gamma_k + \Delta \varphi_{i,t} \quad (1.37)$$

The first thing to note from Equation (1.37) is that the levels of skill prices do not appear; rather only changes are identified (Cortes, 2016, normalizes skill prices to zero in 1976). However, the parameters are still fundamentally non-identified without a further restriction since $\Delta \pi_{k,t}$ and Γ_k are perfectly multicollinear. Either a base period where skill prices do not change (i.e., $\Delta \pi_{k,t} = 0, \forall k$ where $t \in \{1, \dots, T_{base}\}$) is needed, as we do in this paper, or a restriction on Γ_k implicitly made. Stata does this automatically when we implement Estimation (1.33) without base period in the Monte Carlo simulations of Section 1.C.6 below, omitting the skill accumulation parameter for one of the age groups (i.e., setting it to zero). We prefer explicitly defining a base period instead.

A generalized skill accumulation specification: Suppose we have used a base period where skill prices are indeed constant. Does Estimation (1.33) then identify the parameters in the analysis period? Our evidence strongly suggests that experience is occupation-specific. The skill accumulation estimates in Table 1.19 indicate this but also the fact that large wage differences between entrants and incumbents persist when controlling for general age or experience (Figure 1.10). Section 1.B.2 directly rejects the one-dimensional model in favor of multi-dimensional skills (changes). A model that is aligned with the evidence hence allows for this, that is, for example allows for the fact that previous managerial experience imparts more managerial skills than previous experience in production jobs. Equation (1.35) becomes $s_{k,i,t} = \eta_{k,i} + \sum_{k'=1}^K exp_{k',i,t} \cdot \Gamma_{k',k}$, where $exp_{k',i,t} \equiv \sum_{\tau=t_{i,0}}^{t-1} I_{k',i,\tau}$ is the worker's occupation k' specific experience. Running regression (1.33) gives an error term $\varphi_{i,t} = \sum_{k=1}^K I_{k,i,t} [\sum_{k'=1}^K exp_{k',i,t} \cdot \Gamma_{k',k} - (t - t_{i,0}) \cdot \Gamma_k]$ in that case which varies with $I_{k,i,t}$ and is thus systematically related to

⁵⁷We could make the same argument absorbing the fixed effects, i.e., $w_{k,i,t} - \bar{w}_{k,i} = \pi_{k,t} - \bar{\pi}_k + (X_{i,t-1} - \bar{X}_i) \Gamma_k + \varphi_{i,t} - \bar{\varphi}_i$, but Equation (1.37) seems even clearer.

the regressors. This yields biased estimates even without any unobserved idiosyncratic skill shocks that lead to endogenous switching or staying in sectors.

The correct fixed effects regression for skill prices is instead

$$w_{i,t} = \sum_{k=1}^K I_{k,i,t} \pi_{k,t} + \sum_{k=1}^K I_{k,i,t} \eta_{k,i} + \sum_{k=1}^K I_{k,i,t} \sum_{k'=1}^K \exp_{k',i,t} \cdot \Gamma_{k',k} + \varphi_{i,t}, \quad (1.38)$$

that is, it controls for all previous occupation-specific experience separately. While this is conceptually possible to do, its practical implementation is difficult. It introduces many parameters to be estimated (even more when we realistically allow for occupation-specific skill accumulation to vary with age; e.g., see general skill accumulation Equation (1.7) and evidence in Figure 1.6) and it requires high-quality panel data in order to compute the full occupation- and age-specific work experience history of each individual. Cortes (2016) accounts for the fact that labor market experience is occupation-specific by introducing controls for occupation-specific tenure into regression Equation (1.33). In order to deal with the growth in the number of parameters and the length of the employment history that is required for this approach, he assumes that tenure only affects the current job and that workers lose all of its effect once they switch.

We think that occupation-specific tenure is an especially powerful control when at the same time adding separate individual fixed effects for each occupation stint.⁵⁸ That is, to use $\eta_{\lambda(k,i)}$ which differs flexibly for each continuous period $\lambda(k,i) = 1, \dots, \Lambda(k,i)$ in i 's career during which he works in occupation k . Skill accumulation can then be only occupation-specific and straightforwardly interacted with observable characteristics such as age or education (again omitted for brevity):

$$w_{i,t} = \sum_{k=1}^K I_{k,i,t} \pi_{k,t} + \sum_{k=1}^K I_{k,i,t} \eta_{\lambda(k,i)} + \sum_{k=1}^K I_{k,i,t} \cdot \exp_{\lambda(k,i),t} \cdot \Gamma_k + \varphi_{i,t}. \quad (1.39)$$

Here $\exp_{\lambda(k,i),t}$ is the number of years the individual has spent in this occupation stint at time t , which is effectively tenure (but interacted with age, which is omitted in Equation (1.39)) conditional on occupation-stint-specific fixed effects. This is in our view the best specification and close to Cortes (2016)'s arguably most flexible estimation specification. Cavaglia and Etheridge (2017) also employ this specification throughout their analysis and we use it for our empirical robustness checks in the SIAB as well as the Monte Carlo simulations below.

1.B.4.2 Idiosyncratic Skill Shocks

A substantive difference between our proposed method and the fixed effects approach arises in the presence of idiosyncratic skill shocks and endogenous choice, which are indicated by the cross-accumulation parameters of switchers (Table 1.19) as well as higher-than-average skill shocks of occupation incumbents and stayers (Section 1.4.2). We use a simplified analytical argument here.

With idiosyncratic skill shocks, the right-hand-side of skill change Equation (1.35) becomes:

$$s_{k,i,t} = \eta_{k,i} + \sum_{k'=1}^K \sum_{\tau=t_{i,0}}^{t-1} I_{k',i,\tau} \cdot \Gamma_{k',k} + \sum_{\tau=t_{i,0}+1}^t u_{k,i,\tau}. \quad (1.40)$$

⁵⁸In our data, 21% of workers have multiple stints in an occupation during their career.

The regression error in Equation (1.39), $\varphi_{i,t} \equiv \sum_{k=1}^K I_{k,i,t} \sum_{\tau=t_i,0+1}^t u_{k,i,\tau}$, now systematically depends on the full history of previous idiosyncratic skill shocks, which influence current choices (i.e., the regressors in Equation (1.39)). One might expect that the occupation-stint-specific controls in regression (1.39) in principle address this problem, similar to our differenced approach (1.8). But this is not the case.

To see the argument and the bias most clearly suppose for simplicity that all time-varying skill parameters are zero ($\Gamma_{k',k} = 0, \forall k', k$). Suppose also that there are only two sectors, k and a reference occupation k' , and consider first the base period where we assume that $\pi_{k,t} = \pi_{k',t} = \text{const}$ for $t = 1, \dots, T_{\text{base}}$. In this case, simplified wage Equation (1.39) becomes

$$w_{i,t} = \eta_{\lambda(k',i)} + I_{k,i,t} \tilde{\eta}_{\lambda(k,i)} + u_{\lambda(k',i),t} + I_{k,i,t} \tilde{u}_{\lambda(k,i),t} \text{ for } t = 1, \dots, T_{\text{base}}, \quad (1.41)$$

where $\tilde{\eta}_{k,i} \equiv \eta_{\lambda(k,i)} - \eta_{\lambda(k',i)}$ and $\tilde{u}_{k,i,t} \equiv u_{\lambda(k,i),t} - u_{\lambda(k',i),t}$ are relative skill endowments and skill shocks in occupation k . The regression (1.41) is classically endogeneity-biased because the error term $I_{k,i,t} \tilde{u}_{\lambda(k,i),t}$ most likely positively correlates with the regressor $I_{k,i,t} \tilde{\eta}_{\lambda(k,i)}$, even for the stayers in an occupation stint which we are identifying from. This will lead to an overestimation of $\tilde{\eta}_{\lambda(k,i)}$.

If, in order to account for this correlation along the lines of the main text, we introduce saturated choice specific controls, estimation Equation (1.41) becomes

$$w_{i,t} = \eta_{\lambda(k',i)} + I_{k,i,t} \tilde{\eta}_{\lambda(k,i)} + I_{k,i,t} E(\tilde{u}_{\lambda(k,i),t} | I_{k,i,t}, I_{k,i,t-1}) + \nu_{i,t},$$

with $\nu_{i,t} \equiv u_{\lambda(k',i),t} + I_{k,i,t} [\tilde{u}_{\lambda(k,i),t} - E(\tilde{u}_{\lambda(k,i),t} | I_{k,i,t}, I_{k,i,t-1})]$. Since we identify from the wage growth of occupation-stint stayers (for a switcher the fixed effects together with choice specific controls are not identified), consider $E(\tilde{u}_{\lambda(k,i),t} | I_{k,i,t} = I_{k,i,t-1} = 1, \tilde{\eta}_{\lambda(k,i),t}) = E(\tilde{u}_{\lambda(k,i),t} | \tilde{\eta}_{\lambda(k,i),t} > 0, \tilde{\eta}_{\lambda(k,i),t} + \tilde{u}_{\lambda(k,i),t} > 0)$ for occupation k . So even if the correct average $E(\tilde{u}_{\lambda(k,i),t} | I_{k,i,t}, I_{k,i,t-1})$ (not conditioned on $\tilde{\eta}_{\lambda(k,i)}$) were identified, the error term in this regression varies systematically with the fixed effect in the regressor (most likely $\frac{\partial E(\tilde{u}_{\lambda(k,i),t} | \tilde{\eta}_{\lambda(k,i),t} > 0, \tilde{\eta}_{\lambda(k,i),t} + \tilde{u}_{\lambda(k,i),t} > 0)}{\partial \tilde{\eta}_{\lambda(k,i)}} < 0$). Therefore, $\tilde{\eta}_{\lambda(k,i),t}$ identified from occupation k stayers should be downward-biased and also the Γ_k in the full estimation Equation (1.39) are not correctly estimated from the outset. With these (individual-specific) biases from the base period, not only may also the skill price estimates in the analysis period $\pi_{k,t}, t > T_{\text{base}}$ be biased but also will it be hard to sign the direction of these biases. This is a reflection of the fact that fixed effects estimations fundamentally require the exogenous mobility assumption. Relying in the estimation on the wage growth of stayers in an occupation does not in principle alleviate the resulting biases (see Section 1.4.2, which finds that stayers are strongly self-selected according to idiosyncratic skill shocks). The Monte Carlo simulations in Section 1.C.6 show, however, that these biases become quantitatively important only with a lot of switching (i.e., when the dispersion of skill shocks is large).

1.C Monte Carlo Evidence

In this section we provide Monte Carlo evidence for the performance of our and other estimation methods under various assumptions about the data generating process. We first describe the data generating process, which we attempt to keep reasonably close to our sample on the one hand while allowing us to evaluate the impact of key changes on the other hand. We also highlight some stylized facts in the data—for example on occupational switching and the distribution of period-by-period wage changes—that

may help in judging what constitutes reasonable parameter values. We then discuss the actual simulation results in detail.

1.C.1 Data Generating Process

We generate panel datasets similar in structure to the actual SIAB data and set the data generating processes' parameters to the values we will eventually estimate. This allows us to keep some features disciplined by the data while varying components that appear critical to the model. We believe this is much more transparent to the reader than picking arbitrary distributions in a fully stylized setting.

We randomly draw the initial observations of 50,000 individuals from the SIAB as described in Section 1.A.1. The variables we use include initial wages, occupational choices, age (25–54) and year (1975–2010). We then use the parameters from our baseline estimation in Section 1.4 to decompose the initial wage of a worker into a price and a skill component:

$$s_{k,i,t_{i,0}} = w_{i,t_{i,0}} - \pi_{k,t_{i,0}} \quad \text{if } i \in k, \quad (1.42)$$

where $t_{i,0} \in \{1975, \dots, 2010\}$ is the year a worker is first observed.

As the Roy model implies that the initial choice must be optimal in the sense that $w_{k,i,t_{i,0}} \geq w_{k',i,t_{i,0}} \forall k'$, we have a natural bound for the skills a worker possesses in the remaining sectors.

$$s_{k',i,t_{i,0}} \leq s_{k,i,t_{i,0}} + \pi_{k,t_{i,0}} - \pi_{k',t_{i,0}} \quad \text{if } i \in k. \quad (1.43)$$

We draw the initial skills separately and independently for every worker in the sample from a truncated normal distribution with the upper bound given by $s_{k,i,t_{i,0}} + \pi_{k,t_{i,0}} - \pi_{k',t_{i,0}}$. We set the location parameter $\mu_{k,i,t_{i,0}} = s_{k,i,t_{i,0}} + \pi_{k,t_{i,0}}$ and fix the scaling parameter $\sigma = 3$ across workers.

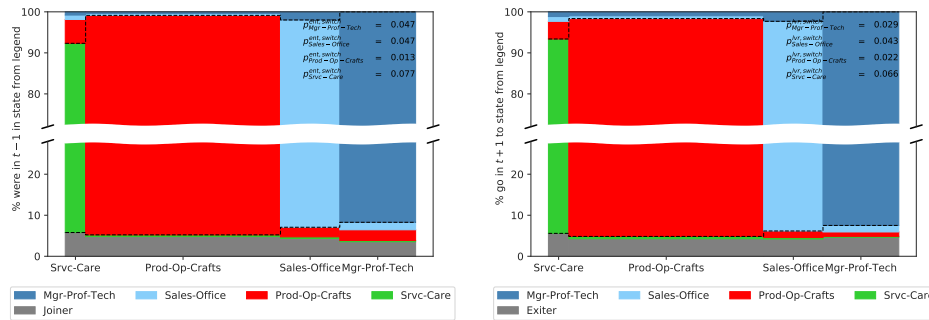
For the following years of a worker's career, we then simulate wage growth as the sum of systematic skill growth and price growth given by $\hat{\Gamma}_{k',k}, \Delta \hat{\pi}_{k,t}$. On top of that, we add idiosyncratic skill shocks depending on the specification and finally let workers choose their preferred sector based on comparative advantage (possibly including costs of switching and changing non monetary amenities). We repeat this until a worker's maximum age of 54 is reached or the sample period ends. We rerun the exercise for 100 Monte Carlo repetitions and estimate price and skill changes on each sample. We then compute the average price trends and skill accumulation function across repetitions. For computational reasons, we only use four occupations (i.e., the four broad occupation groups of the main text). In the subsections from 1.C.3 onward, we report the results by comparing estimated and true parameter values $\hat{\pi}_{k,t}, \pi_{k,t}$ for different Monte Carlo specifications.

1.C.2 Key facts in the SIAB data


We first document some facts in the SIAB data. While the goal of our Monte Carlo studies is not to replicate these facts it is useful to keep them in mind in order to judge what may be reasonable parameter values. For example, if a specification yields consistent parameter estimates but switching between occupations is shut down completely (as would happen, for example, if switching costs are high and idiosyncratic shocks are small), this would not constitute a very realistic exercise.

Figure 1.16: Descriptive statistics in the SIAB data

(a) Occupation entrants/incumb. in $t + 1$ (b) Occupation leavers/stayers in $t - 1$



(c) Distribution of annual wage growth (d) Evolution of the wage distribution



Notes: The amount of stayers in the switching graphs corresponds to the area between the two dashed lines. The occupations are ordered by the average wage within the occupations across years 1985 until 2010. Entering or leaving an occupation to an occupation with a higher average wage is depicted by an area above the upper dashed line and vice versa for occupations with a lower average wage. The wage growth histogram was calculated using 100 equally sized bins between -0.5 and 0.5. Panels (a) and (b) use the specification with filled-up unemployment and out-of-the-labor-force spells (see Section 1.A.1.4) for better comparability later on, since unemployment or out-of-the-labor-force do not form a part of our Monte Carlo exercise.

Figure 1.16a shows the composition of incumbents and entrants into occupation groups from t to $t + 1$. The groups are ranked horizontally by their average wages over the years 1985–2010. They are ranked vertically also by average wages, with the incumbents in the middle and lower/higher earning occupations from which workers enter bottom/top, respectively. Figure 1.16b depicts the same thing for stayers and leavers from $t - 1$ to t . Entrants into our sample at age 25 (joiners) and leavers at age 54 (exiters) are coded in gray at the very bottom. Figures 1.16a and 1.16b use the specification with filled-up unemployment and out-of-the-labor-force spells (see Section 1.A.1.4) for better comparability later on, since unemployment or out-of-the-labor-force do not form a part of our Monte Carlo exercise.

Both panels of Figure 1.16 show that a substantial amount of occupation entry and exit are at the ends of the sample age range. But there is also a large share of Srvc-Care workers who switch to Prod-Op-Crafts and even more who switch from Prod-Op-Crafts

into Svc-Care, which is consistent with a growth of Svc-Card and decline of Prod-Op-Crafts over time. There are in addition substantial switches between Sales-Office and Mgr-Prof-Tech, among others.

The bottom panels of Figure 1.16 depict a histogram of workers' annual wage growth and repeat the distribution of key quantiles of the wage distribution for comparability.

1.C.3 Baseline model

This subsection shows the results for a Monte Carlo exercise where wage growth only stems from price growth or systematic skill growth but not from idiosyncratic shocks. The primary reason for this specification is to assess the quality of the approximation in Equation (1.4).

Despite the absence of shocks, Figure 1.17 shows that there are a substantial number of switchers just because of prices and (cross-) skill accumulation. Unsurprisingly, other statistics do not match up well (for example, the standard deviation of annual wage growth is smaller by a factor of 7 than in Figure 1.16c), but that is not the point in this particular exercise.

Figure 1.18 shows cumulative skill prices and skills in the four occupations, with estimates depicted in the solid lines and true parameter values as crosses. Clearly, the proposed method is able to estimate skill price and skills trends from observed wage changes. Note this is the case for all 100 experiments as no individual lines for specific experiments are to be seen. The only randomness in these cases comes from the initial distributions and it should be taken care of by our estimates if the approximation in Equation (1.4) works well. The graph certainly suggests that it does.

Idiosyncratic skill shocks $u_{k,i,t}$ as in Equation (1.7) introduce an endogeneity bias to our estimates, which is rising with the standard deviation of the shocks. Figure 1.19 shows information about the amount of switchers and the dispersion of log wage growth under a scenario where the dispersion of (normally distributed, center 0) skill shocks is $\sigma = 0.5 \cdot \sigma_{\Delta \log(w_{i,t})}^{SIAB}$. Clearly, both the distribution of wage growth as well as the distribution of switchers is more comparable in Monte Carlo and SIAB data in this setting.

Figure 1.20 shows that, as predicted in Section 1.3.3, the OLS estimates contain modest upward bias of stayers' skill accumulation coefficients, whereas the IV estimates are almost exactly on target. Also as expected, the cross-accumulation parameters are generally upward-biased in the OLS (Table 1.7); and in the IV with weak instruments, they are large in absolute values (Table 1.8). Both sets of skill price estimates track the evolution of their actual values very closely on average (individual experiments in the shimmering lines do modestly vary around the truth).

Next, we increase the standard deviation of $u_{k,i,t}$ to 1.5 times the standard deviation of log wage growth observed in the SIAB data. Figure 1.21 shows that the number of switches in the Monte Carlo sample increase a lot and it is now much larger than in the actual SIAB data. Additionally, the simulated distribution of log wage growth and overall wage inequality are much more dispersed compared to the observed distributions. So this is clearly an extreme setting, which we create to see whether the bias due to skill shocks can in some instances become substantial. The OLS estimates in Figure 1.22 show that the bias in the skill accumulation parameters indeed becomes large. However, skill price changes are only slightly downward biased and not far off their targets (the bias is a bit larger for Mgr-Prof-Tech).

In Panels 1.22c and 1.22d, we implement the instrumental variables strategy that was outlined in the main text in order to deal with the remaining bias due to idiosyncratic skill shocks. It turns out that IV does indeed estimate the correct skill price changes even in this extreme case. The skill accumulation parameters are still upward biased, but less so than in the OLS. Therefore, the instrumental variables estimation is robust even to rather extreme idiosyncratic skill shocks. It seems that implementing the IV will be helpful in practice to check whether the (main) OLS estimates might be biased because of such large shocks.

Finally, we make skill shocks persistent by introducing autocorrelation between shocks in t and $t-x$ as $u_{k,i,t} = 0.3 \cdot u_{k,i,t-1} + \eta_{k,i,t}$. We calibrate $\sigma = 0.5 \cdot \sigma_{\Delta \log(w_{i,t})}^{SIAB}$.⁵⁹ When workers switch because of a positive skill shocks, they move into an occupation where previous skill shocks tend to be positive as well due to the autocorrelation coefficient ρ . The results are displayed in Figure 1.24. Again, any bias in the skill price estimates for both the OLS and IV is minor, while the structural skill accumulation parameters are unsurprisingly quite off.

1.C.3.1 No Idiosyncratic Skill Shocks


Table 1.5: Parameters

Parameter	Value
N	50000
Repetitions	100
Skill shocks in k	uniform, $(\mu, \sigma_k) = (0, 0 \cdot \sigma_{\Delta \log(w_i)}^{SIAB})$
Stayers accumulation $\gamma_{k,k,a}, k' = k$	$\hat{\gamma}_{k,k,a}^{SIAB}$
Cross accumulation $\gamma_{k',k,a}, k' \neq k$	$\frac{1}{3} \hat{\gamma}_{k',k,a}^{SIAB}$
ρ in $\varepsilon_{i,t} = \rho \varepsilon_{i,t-1} + v_{i,t}$	0
Switching costs c	0
Amenity trends, $t = 1985, \dots, 2010$	$[\Delta \Psi_{k,t}]_{k=1,\dots,4} = [0, 0, 0, 0]$


⁵⁹The evidence on the existence and importance of correlated wage shocks seems to be mixed (e.g., see Gibbons and Waldman, 1999, and the references therein).

Figure 1.17: Descriptives, no shocks


(a) Occupation entrants/incumb. in $t + 1$ (b) Occupation leavers/stayers in $t - 1$



(c) Distribution of annual wage growth

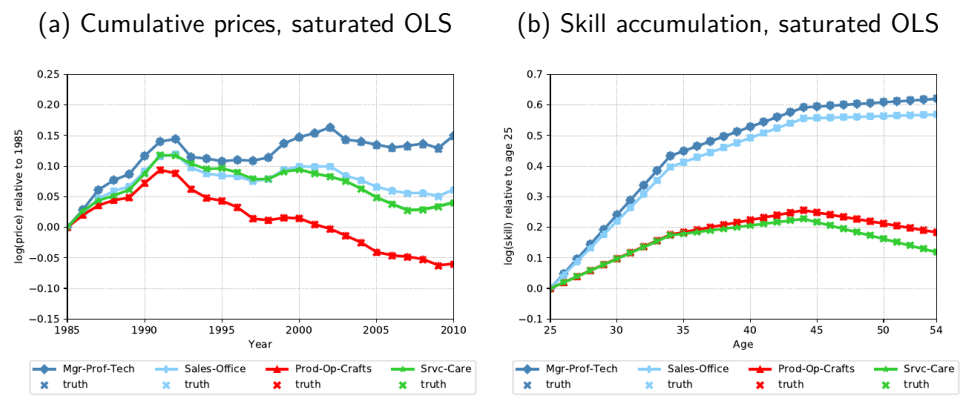


(d) Evolution of the wage distribution



Notes: The amount of stayers in the switching graphs corresponds to the area between the two dashed lines. The occupations are ordered by the average wage within the occupations across years 1985 until 2010. Entering or leaving an occupation to an occupation with a higher average wage is depicted by an area above the upper dashed line and vice versa for occupations with a lower average wage. The wage growth histogram was calculated using 100 equally sized bins between -0.5 and 0.5. The results are averages across the 100 Monte Carlo replications.

Figure 1.18: Estimation results, no shocks



Notes: Crosses “x” represent true values used to simulate workers’ wages. Solid lines are averages across the 100 Monte Carlo replications. Shimmering lines in the background represent individual Monte Carlo replications. There is no variation in the background to be seen as there is no randomness included. OLS estimates as described by Equation (1.9).


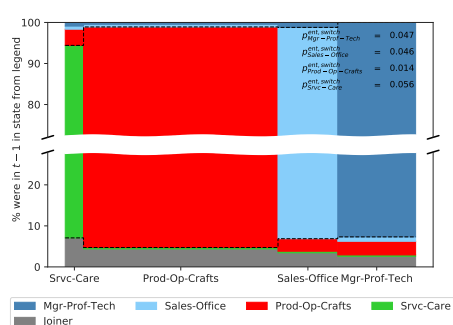
1.C.3.2 Moderately Dispersed Shocks

Table 1.6: Parameters


Parameter	Value
N	50000
Repetitions	100
Skill shocks in k	observed wage growth distribution, $(\mu, \sigma_k) = (0, 0.5 \cdot \sigma_{\Delta \log(w_i)}^{SIAB})$
Stayers accumulation $\gamma_{k,k,a}, k' = k$	$\hat{\gamma}_{k,k,a}^{SIAB}$
Cross accumulation $\gamma_{k',k,a}, k' \neq k$	$\frac{1}{3} \hat{\gamma}_{k',k,a}^{SIAB}$
ρ in $\varepsilon_{i,t} = \rho \varepsilon_{i,t-1} + v_{i,t}$	0
Switching costs c	0
Amenity trends, $t = 1985, \dots, 2010$	$[\Delta \Psi_{k,t}]_{k=1,\dots,4} = [0, 0, 0, 0]$

Figure 1.19: Descriptives, moderate shocks


(a) Occupation entrants/incumb. in $t + 1$ (b) Occupation leavers/stayers in $t - 1$



(c) Distribution of annual wage growth

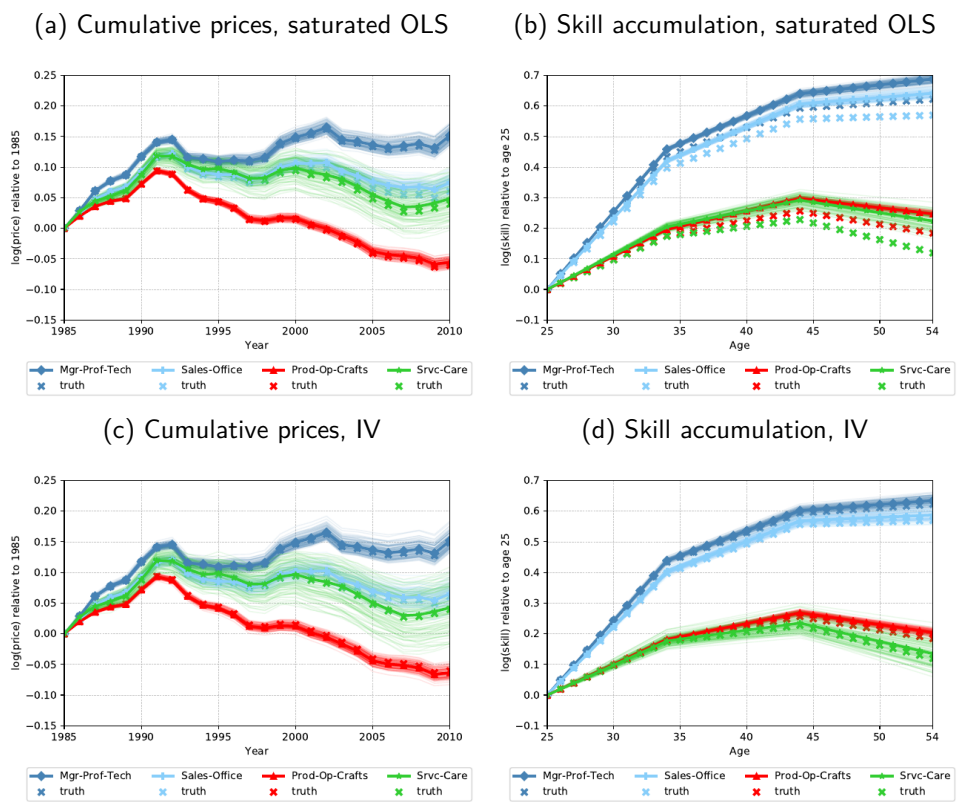


(d) Evolution of the wage distribution



Notes: The amount of stayers in the switching graphs corresponds to the area between the two dashed lines. The occupations are ordered by the average wage within the occupations across years 1985 until 2010. Entering or leaving an occupation to an occupation with a higher average wage is depicted by an area above the upper dashed line and vice versa for occupations with a lower average wage. The wage growth histogram was calculated using 100 equally sized bins between -0.5 and 0.5. The results are averages across the 100 Monte Carlo replications. OLS estimates as described by Equation (1.9). IV estimates as described at the end of Section 1.3.3.

Figure 1.20: Estimation results, moderate shocks



Notes: Crosses “x” represent true values used to simulate workers’ wages. Solid lines are averages across the 100 Monte Carlo replications. Shimmering lines in the background represent individual Monte Carlo replications. OLS estimates as described by Equation (1.9). IV estimates as described at the end of Section 1.3.3.

Table 1.7: True and estimated skill accumulation parameters, saturated OLS

Previous sector	Current sector		Age group					
			[25, 34]		[35, 44]		[45, 54]	
			$\hat{\gamma}_{k',k,a}$	$\gamma_{k',k,a}^{true}$	$\hat{\gamma}_{k',k,a}$	$\gamma_{k',k,a}^{true}$	$\hat{\gamma}_{k',k,a}$	$\gamma_{k',k,a}^{true}$
Mgr-Prof-Tech	Mgr-Prof-Tech	γ	0.051	0.048	0.018	0.016	0.005	0.003
		σ_γ	0.000		0.000		0.000	
	Sales-Office	γ	0.069	0.063	0.002	0.009	-0.045	-0.010
		σ_γ	0.009		0.012		0.021	
	Prod-Op-Crafts	γ	0.031	0.023	-0.032	-0.011	-0.050	-0.022
		σ_γ	0.014		0.014		0.017	
	Srvc-Care	γ	-0.093	-0.008	-0.131	-0.036	-0.031	-0.004
		σ_γ	0.061		0.069		0.034	
Sales-Office	Mgr-Prof-Tech	γ	0.098	0.088	0.030	0.027	0.003	0.009
		σ_γ	0.006		0.009		0.013	
	Sales-Office	γ	0.047	0.044	0.019	0.016	0.004	0.001
		σ_γ	0.000		0.000		0.000	
	Prod-Op-Crafts	γ	0.070	0.056	0.032	0.019	-0.008	-0.008
		σ_γ	0.007		0.007		0.012	
	Srvc-Care	γ	-0.042	0.010	-0.125	-0.034	-0.083	-0.024
		σ_γ	0.042		0.067		0.065	
Prod-Op-Crafts	Mgr-Prof-Tech	γ	0.086	0.075	0.035	0.042	0.012	0.021
		σ_γ	0.004		0.005		0.006	
	Sales-Office	γ	0.037	0.036	0.017	0.022	-0.014	0.000
		σ_γ	0.006		0.006		0.009	
	Prod-Op-Crafts	γ	0.022	0.020	0.010	0.008	-0.005	-0.007
		σ_γ	0.000		0.000		0.000	
	Srvc-Care	γ	-0.056	-0.017	-0.076	-0.014	-0.038	-0.009
		σ_γ	0.022		0.025		0.016	
Srvc-Care	Mgr-Prof-Tech	γ	0.107	0.099	0.063	0.063	0.047	0.041
		σ_γ	0.011		0.014		0.014	
	Sales-Office	γ	0.092	0.090	0.048	0.048	0.014	0.015
		σ_γ	0.012		0.016		0.026	
	Prod-Op-Crafts	γ	0.119	0.106	0.085	0.075	0.055	0.037
		σ_γ	0.007		0.008		0.008	
	Srvc-Care	γ	0.023	0.019	0.009	0.005	-0.007	-0.011
		σ_γ	0.001		0.001		0.001	

Notes: The table shows the estimated $\hat{\gamma}_{k',k,a}$, which is a scalar element of $\Gamma_{k',k}$ representing skill accumulation of age group a . The four groups are based on an aggregation of detailed occupations in the SIAB SUF as described in Appendix Table 1.4. OLS estimates as described by Equation (1.9).

Table 1.8: True and estimated skill accumulation parameters, IV

Previous sector	Current sector		Age group					
			[25, 34]		[35, 44]		[45, 54]	
			$\hat{\gamma}_{k',k,a}$	$\gamma_{k',k,a}^{true}$	$\hat{\gamma}_{k',k,a}$	$\gamma_{k',k,a}^{true}$	$\hat{\gamma}_{k',k,a}$	$\gamma_{k',k,a}^{true}$
Mgr-Prof-Tech	Mgr-Prof-Tech	γ	0.049	0.048	0.016	0.016	0.003	0.003
		σ_γ	0.001		0.000		0.000	
	Sales-Office	γ	0.173	0.063	0.162	0.009	0.152	-0.010
		σ_γ	0.016		0.024		0.038	
	Prod-Op-Crafts	γ	0.249	0.023	0.210	-0.011	0.192	-0.022
		σ_γ	0.030		0.027		0.030	
	Srvc-Care	γ	0.283	-0.008	0.253	-0.036	0.152	-0.004
		σ_γ	0.141		0.134		0.054	
Sales-Office	Mgr-Prof-Tech	γ	0.164	0.088	0.132	0.027	0.123	0.009
		σ_γ	0.010		0.017		0.029	
	Sales-Office	γ	0.045	0.044	0.017	0.016	0.002	0.001
		σ_γ	0.001		0.001		0.001	
	Prod-Op-Crafts	γ	0.171	0.056	0.152	0.019	0.154	-0.008
		σ_γ	0.014		0.015		0.025	
	Srvc-Care	γ	0.228	0.010	0.237	-0.034	0.169	-0.024
		σ_γ	0.095		0.134		0.109	
Prod-Op-Crafts	Mgr-Prof-Tech	γ	0.150	0.075	0.125	0.042	0.105	0.021
		σ_γ	0.013		0.013		0.015	
	Sales-Office	γ	0.133	0.036	0.122	0.022	0.109	0.000
		σ_γ	0.013		0.011		0.019	
	Prod-Op-Crafts	γ	0.020	0.020	0.009	0.008	-0.006	-0.007
		σ_γ	0.000		0.000		0.000	
	Srvc-Care	γ	0.222	-0.017	0.180	-0.014	0.119	-0.009
		σ_γ	0.071		0.049		0.028	
Srvc-Care	Mgr-Prof-Tech	γ	0.171	0.099	0.134	0.063	0.108	0.041
		σ_γ	0.041		0.041		0.025	
	Sales-Office	γ	0.159	0.090	0.124	0.048	0.100	0.015
		σ_γ	0.029		0.047		0.061	
	Prod-Op-Crafts	γ	0.180	0.106	0.147	0.075	0.113	0.037
		σ_γ	0.018		0.018		0.016	
	Srvc-Care	γ	0.020	0.019	0.006	0.005	-0.010	-0.011
		σ_γ	0.002		0.002		0.001	

Notes: The table shows the estimated $\hat{\gamma}_{k',k,a}$, which is a scalar element of $\Gamma_{k',k}$ representing skill accumulation of age group a . The four groups are based on an aggregation of detailed occupations in the SIAB SUF as described in Appendix Table 1.4. IV estimates as described at the end of Section 1.3.3.


1.C.3.3 Highly Dispersed Shocks

Table 1.9: Parameters


Parameter	Value
N	50000
Repetitions	100
Skill shocks in k	observed wage growth distribution, $(\mu, \sigma_k) = (0, 1.5 \cdot \sigma_{\Delta \log(w_i)}^{SIAB})$
Stayers accumulation $\gamma_{k,k,a}, k' = k$	$\hat{\gamma}_{k,k,a}^{SIAB}$
Cross accumulation $\gamma_{k',k,a}, k' \neq k$	$\frac{1}{3} \hat{\gamma}_{k',k,a}^{SIAB}$
ρ in $\varepsilon_{i,t} = \rho \varepsilon_{i,t-1} + v_{i,t}$	0
Switching costs c	0
Amenity trends, $t = 1985, \dots, 2010$	$[\Delta \Psi_{k,t}]_{k=1,\dots,4} = [0, 0, 0, 0]$

Figure 1.21: Descriptives, highly dispersed shocks

(a) Occupation entrants/incumb. in $t + 1$ (b) Occupation leavers/stayers in $t - 1$

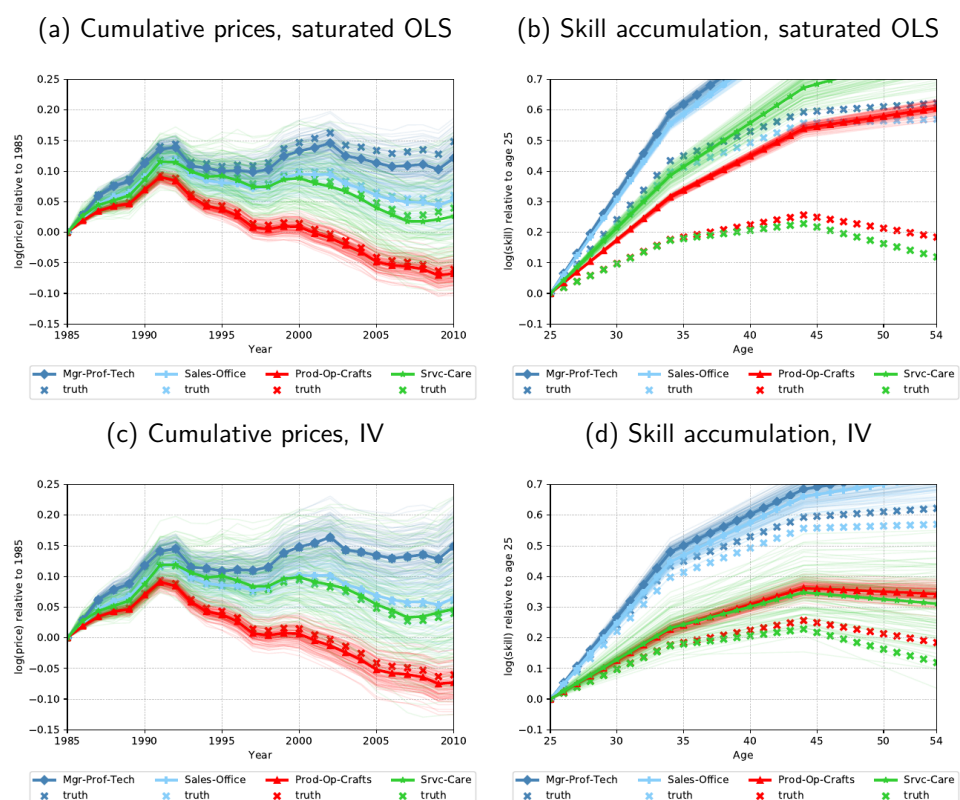


(c) Distribution of annual wage growth (d) Evolution of the wage distribution



Notes: The amount of stayers in the switching graphs corresponds to the area between the two dashed lines. The occupations are ordered by the average wage within the occupations across years 1985 until 2010. Entering or leaving an occupation to an occupation with a higher average wage is depicted by an area above the upper dashed line and vice versa for occupations with a lower average wage. The wage growth histogram was calculated using 100 equally sized bins between -0.5 and 0.5. The results are averages across the 100 Monte Carlo replications.

Figure 1.22: Estimation results, highly dispersed shocks



Notes: Crosses “x” represent true values used to simulate workers’ wages. Solid lines are averages across the 100 Monte Carlo replications. Shimmering lines in the background represent individual Monte Carlo replications. OLS estimates as described by Equation (1.9). IV estimates as described at the end of Section 1.3.3.

Table 1.10: True and estimated skill accumulation parameters, saturated OLS

Previous sector	Current sector		Age group						
			[25, 34]		[35, 44]		[45, 54]		
			$\hat{\gamma}_{k',k,a}$	$\gamma_{k',k,a}^{true}$	$\hat{\gamma}_{k',k,a}$	$\gamma_{k',k,a}^{true}$	$\hat{\gamma}_{k',k,a}$	$\gamma_{k',k,a}^{true}$	
Mgr-Prof-Tech	Mgr-Prof-Tech	γ	0.065	0.048	0.031	0.016	0.017	0.003	
		σ_γ	0.001		0.001		0.001		
	Sales-Office	γ	0.113	0.063	0.037	0.009	-0.001	-0.010	
		σ_γ	0.022		0.023		0.031		
	Prod-Op-Crafts	γ	0.097	0.023	0.030	-0.011	0.002	-0.022	
		σ_γ	0.020		0.019		0.023		
	Srvc-Care	γ	0.038	-0.008	-0.008	-0.036	0.010	-0.004	
		σ_γ	0.052		0.053		0.043		
	Sales-Office	Mgr-Prof-Tech	γ	0.130	0.088	0.069	0.027	0.038	0.009
			σ_γ	0.018		0.020		0.026	
Sales-Office		γ	0.062	0.044	0.032	0.016	0.016	0.001	
		σ_γ	0.001		0.001		0.001		
Prod-Op-Crafts		γ	0.116	0.056	0.073	0.019	0.032	-0.008	
		σ_γ	0.018		0.017		0.023		
Srvc-Care		γ	0.058	0.010	0.010	-0.034	0.021	-0.024	
		σ_γ	0.044		0.056		0.060		
Prod-Op-Crafts		Mgr-Prof-Tech	γ	0.103	0.075	0.052	0.042	0.026	0.021
			σ_γ	0.011		0.013		0.016	
	Sales-Office	γ	0.062	0.036	0.036	0.022	0.008	0.000	
		σ_γ	0.015		0.016		0.021		
	Prod-Op-Crafts	γ	0.035	0.020	0.022	0.008	0.007	-0.007	
		σ_γ	0.001		0.001		0.001		
	Srvc-Care	γ	0.008	-0.017	0.000	-0.014	0.002	-0.009	
		σ_γ	0.026		0.027		0.026		
	Srvc-Care	Mgr-Prof-Tech	γ	0.154	0.099	0.121	0.063	0.107	0.041
			σ_γ	0.025		0.027		0.028	
Sales-Office		γ	0.143	0.090	0.116	0.048	0.094	0.015	
		σ_γ	0.027		0.033		0.043		
Prod-Op-Crafts		γ	0.159	0.106	0.129	0.075	0.101	0.037	
		σ_γ	0.016		0.016		0.019		
Srvc-Care		γ	0.043	0.019	0.029	0.005	0.012	-0.011	
		σ_γ	0.002		0.002		0.002		

Notes: The table shows the estimated $\hat{\gamma}_{k',k,a}$, which is a scalar element of $\Gamma_{k',k}$, representing skill accumulation of age group a . The four groups are based on an aggregation of detailed occupations in the SIAB SUF as described in Appendix Table 1.4. OLS estimates as described by Equation (1.9).

Table 1.11: True and estimated skill accumulation parameters, IV

Previous sector	Current sector		Age group					
			[25, 34]		[35, 44]		[45, 54]	
			$\hat{\gamma}_{k',k,a}$	$\gamma_{k',k,a}^{true}$	$\hat{\gamma}_{k',k,a}$	$\gamma_{k',k,a}^{true}$	$\hat{\gamma}_{k',k,a}$	$\gamma_{k',k,a}^{true}$
Mgr-Prof-Tech	Mgr-Prof-Tech	γ	0.053	0.048	0.020	0.016	0.007	0.003
		σ_γ	0.002		0.002		0.002	
	Sales-Office	γ	0.406	0.063	0.405	0.009	0.403	-0.010
		σ_γ	0.050		0.052		0.068	
	Prod-Op-Crafts	γ	0.477	0.023	0.442	-0.011	0.437	-0.022
		σ_γ	0.049		0.044		0.051	
	Srvc-Care	γ	0.515	-0.008	0.488	-0.036	0.385	-0.004
		σ_γ	0.143		0.120		0.083	
Sales-Office	Mgr-Prof-Tech	γ	0.378	0.088	0.371	0.027	0.359	0.009
		σ_γ	0.045		0.048		0.062	
	Sales-Office	γ	0.050	0.044	0.021	0.016	0.007	0.001
		σ_γ	0.002		0.002		0.002	
	Prod-Op-Crafts	γ	0.417	0.056	0.396	0.019	0.393	-0.008
		σ_γ	0.043		0.041		0.054	
	Srvc-Care	γ	0.470	0.010	0.462	-0.034	0.397	-0.024
		σ_γ	0.119		0.126		0.130	
Prod-Op-Crafts	Mgr-Prof-Tech	γ	0.353	0.075	0.351	0.042	0.339	0.021
		σ_γ	0.035		0.034		0.041	
	Sales-Office	γ	0.376	0.036	0.368	0.022	0.362	0.000
		σ_γ	0.044		0.039		0.049	
	Prod-Op-Crafts	γ	0.025	0.020	0.013	0.008	-0.002	-0.007
		σ_γ	0.001		0.001		0.001	
	Srvc-Care	γ	0.435	-0.017	0.414	-0.014	0.371	-0.009
		σ_γ	0.076		0.060		0.056	
Srvc-Care	Mgr-Prof-Tech	γ	0.363	0.099	0.356	0.063	0.325	0.041
		σ_γ	0.080		0.078		0.068	
	Sales-Office	γ	0.374	0.090	0.360	0.048	0.339	0.015
		σ_γ	0.076		0.092		0.114	
	Prod-Op-Crafts	γ	0.369	0.106	0.351	0.075	0.324	0.037
		σ_γ	0.045		0.045		0.051	
	Srvc-Care	γ	0.026	0.019	0.012	0.005	-0.004	-0.011
		σ_γ	0.004		0.004		0.004	

Notes: The table shows the estimated $\hat{\gamma}_{k',k,a}$, which is a scalar element of $\Gamma_{k',k}$ representing skill accumulation of age group a . The four groups are based on an aggregation of detailed occupations in the SIAB SUF as described in Appendix Table 1.4. IV estimates as described at the end of Section 1.3.3.

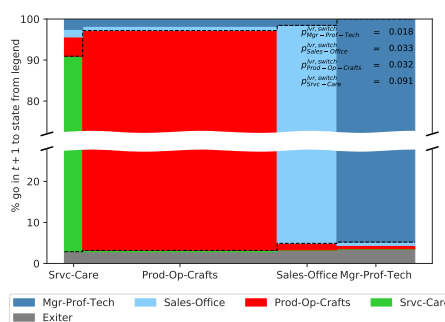
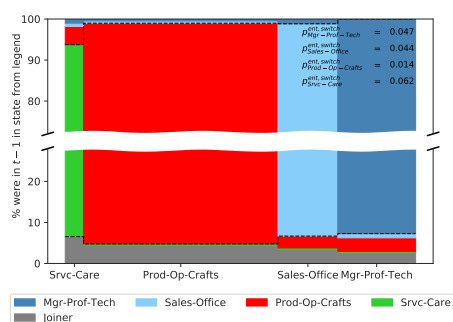
1.C.3.4 Persistent Shocks

Table 1.12: Parameters


Parameter	Value
N	50000
Repetitions	100
Skill shocks in k	observed wage growth distribution, $(\mu, \sigma_k) = (0, 0.5 \cdot \sigma_{\Delta \log(w_i)}^{SIAB})$
Stayers accumulation $\gamma_{k,k,a}, k' = k$	$\hat{\gamma}_{k,k,a}^{SIAB}$
Cross accumulation $\gamma_{k',k,a}, k' \neq k$	$\frac{1}{3} \hat{\gamma}_{k',k,a}^{SIAB}$
ρ in $\varepsilon_{i,t} = \rho \varepsilon_{i,t-1} + v_{i,t}$	0.3
Switching costs c	0
Amenity trends, $t = 1985, \dots, 2010$	$[\Delta \Psi_{k,t}]_{k=1,\dots,4} = [0, 0, 0, 0]$

Figure 1.23: Descriptives, persistent shocks


(a) Occupation entrants/incumb. in $t + 1$ (b) Occupation leavers/stayers in $t - 1$



(c) Distribution of annual wage growth

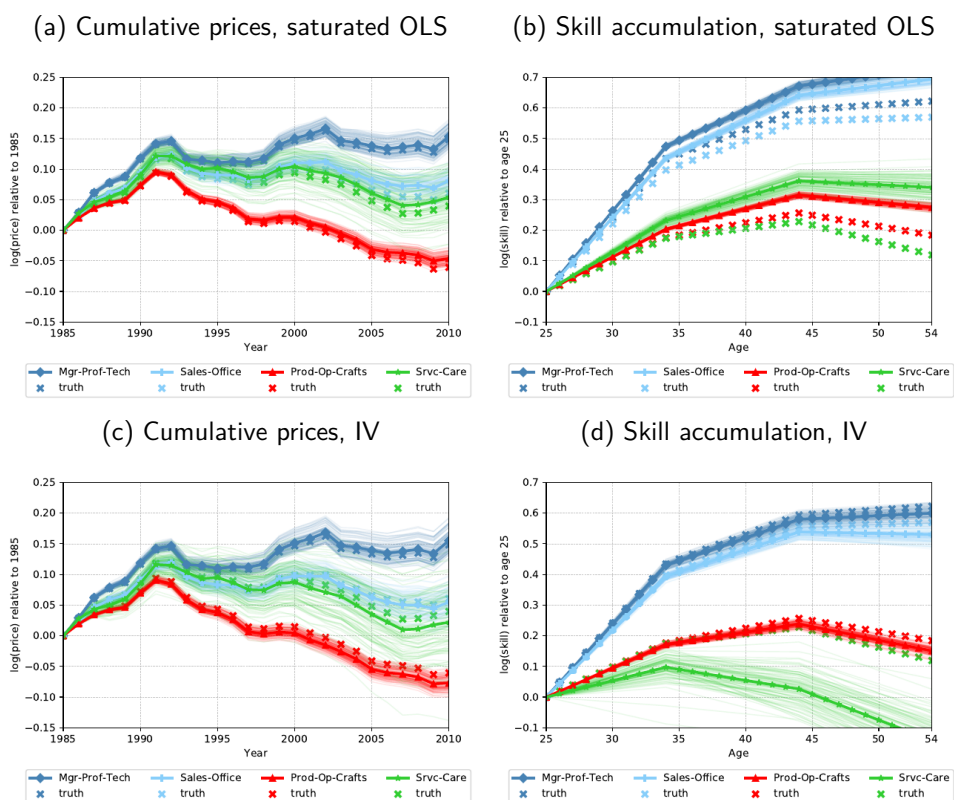


(d) Evolution of the wage distribution



Notes: The amount of stayers in the switching graphs corresponds to the area between the two dashed lines. The occupations are ordered by the average wage within the occupations across years 1985 until 2010. Entering or leaving an occupation to an occupation with a higher average wage is depicted by an area above the upper dashed line and vice versa for occupations with a lower average wage. The wage growth histogram was calculated using 100 equally sized bins between -0.5 and 0.5. The results are averages across the 100 Monte Carlo replications.

Figure 1.24: Estimation results, persistent shocks



Notes: Crosses “x” represent true values used to simulate workers’ wages. Solid lines are averages across the 100 Monte Carlo replications. Shimmering lines in the background represent individual Monte Carlo replications. OLS estimates as described by Equation (1.9). IV estimates as described at the end of Section 1.3.3.

1.C.4 Switching Costs

In this section, we make the decision problem depend on non wage switching costs as a fraction of log wages. We assume that every worker has to pay a (psychic) utility cost when wanting to switch so that the potential utility amounts to $U = w$ when not switching and $U = (1-c)w$ when switching. We start with switching costs of $c = 0.075$. First, in a model without skill shocks and with switching costs, our approximation (1.4) continues to work well (Figures 1.25 and 1.26). We have trebled the size of cross-accumulation parameters in the data generating process here (see Table 1.13), since otherwise we would hardly observe any switches without skill shocks and given the switching costs.

The switching cost also reduces the number of switchers under moderate skill shocks as depicted in Figure 1.27. Figure 1.28 shows that it does not bias our results. In fact, as skill prices are identified well from wage growth of stayers, switching costs make any bias of the skill prices less severe. Both OLS and IV therefore work for this scenario with moderate shocks and moderate switching costs, which we deem a rather realistic one.

An additional reason for why we think that the scenario with moderate skill shocks and moderate switching costs is a sensible benchmark are the estimated (cross-) accumulation parameters. This is a bit subtle: Note that OLS's upward-bias of the cross-accumulation parameters ($\hat{\Gamma}_{k',k}$ for $k' \neq k$) is exacerbated here because even larger idiosyncratic skill shocks are required to overcome the switch costs (Table 1.15). The IV's weak instrument problem for the switchers is also more severe (Table 1.16). The data generating process of the Monte Carlos used a third of the cross-accumulation values estimated in the SIAB (see Table 1.14). The estimates based on the simulations with moderate shocks and moderate switching costs re-create these values, i.e., they overstate the target by a factor of three on average. Therefore, this scenario approximately "replicates its own bias" in the estimation of the cross-accumulation parameters. That is, what we pick as data generating process and the estimates that we receive are consistent with one another.

Finally, we increase the switching costs to $c = 0.2$ and the standard deviation of skill shocks to 1.5 times the standard deviation of log wage growth in the SIAB. Once again, the estimates of the skill prices, and especially in the IV, are quite close to their true values.


1.C.4.1 Benchmark: Moderate Switching Costs, No Shocks

Table 1.13: Parameters


Parameter	Value
N	50000
Repetitions	100
Skill shocks in k	uniform, $(\mu, \sigma_k) = (0, 0 \cdot \sigma_{\Delta \log(w_i)}^{SIAB})$
Stayers accumulation $\gamma_{k,k,a}, k' = k$	$\hat{\gamma}_{k,k,a}^{SIAB}$
Cross accumulation $\gamma_{k',k,a}, k' \neq k$	$\frac{1}{I} \hat{\gamma}_{k',k,a}^{SIAB}$
ρ in $\varepsilon_{i,t} = \rho \varepsilon_{i,t-1} + v_{i,t}$	0
Switching costs c	0.05
Amenity trends, $t = 1985, \dots, 2010$	$[\Delta \Psi_{k,t}]_{k=1,\dots,4} = [0, 0, 0, 0]$

Figure 1.25: Descriptives, moderate switch costs, no shocks

(a) Occupation entrants/incumb. in $t + 1$ (b) Occupation leavers/stayers in $t - 1$

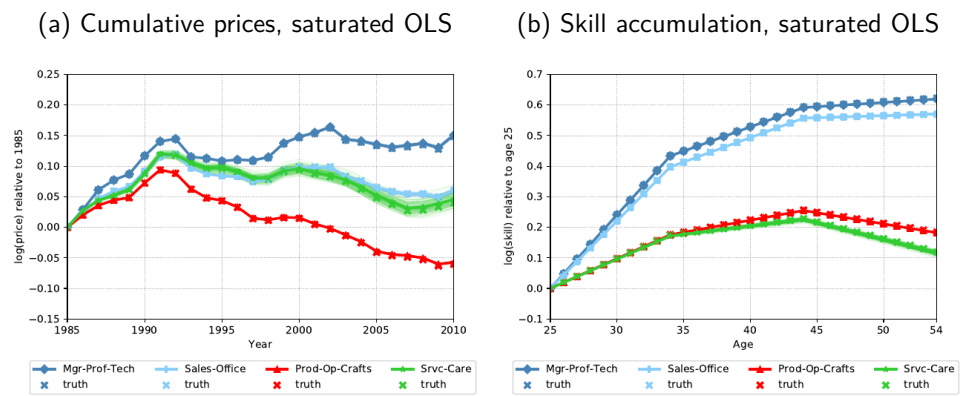


(c) Distribution of annual wage growth (d) Evolution of the wage distribution



Notes: The amount of stayers in the switching graphs corresponds to the area between the two dashed lines. The occupations are ordered by the average wage within the occupations across years 1985 until 2010. Entering or leaving an occupation to an occupation with a higher average wage is depicted by an area above the upper dashed line and vice versa for occupations with a lower average wage. The wage growth histogram was calculated using 100 equally sized bins between -0.5 and 0.5. The results are averages across the 100 Monte Carlo replications.

Figure 1.26: Estimation results, moderate switch costs, no shocks



Notes: Crosses “x” represent true values used to simulate workers’ wages. Solid lines are averages across the 100 Monte Carlo replications. Shimmering lines in the background represent individual Monte Carlo replications. OLS estimates as described by Equation (1.9).



1.C.4.2 Moderate Switching Costs, Moderate Shocks

Table 1.14: Parameters


Parameter	Value
N	50000
Repetitions	100
Skill shocks in k	observed wage growth distribution, $(\mu, \sigma_k) = (0, 0.5 \cdot \sigma_{\Delta \log(w_i)}^{SIAB})$
Stayers accumulation $\gamma_{k,k,a}, k' = k$	$\hat{\gamma}_{k,k,a}^{SIAB}$
Cross accumulation $\gamma_{k',k,a}, k' \neq k$	$\frac{1}{3} \hat{\gamma}_{k',k,a}^{SIAB}$
ρ in $\varepsilon_{i,t} = \rho \varepsilon_{i,t-1} + v_{i,t}$	0
Switching costs c	0.05
Amenity trends, $t = 1985, \dots, 2010$	$[\Delta \Psi_{k,t}]_{k=1,\dots,4} = [0, 0, 0, 0]$

Figure 1.27: Descriptives, moderate switching costs and moderate shocks


(a) Occupation entrants/incumb. in $t + 1$ (b) Occupation leavers/stayers in $t - 1$



(c) Distribution of annual wage growth

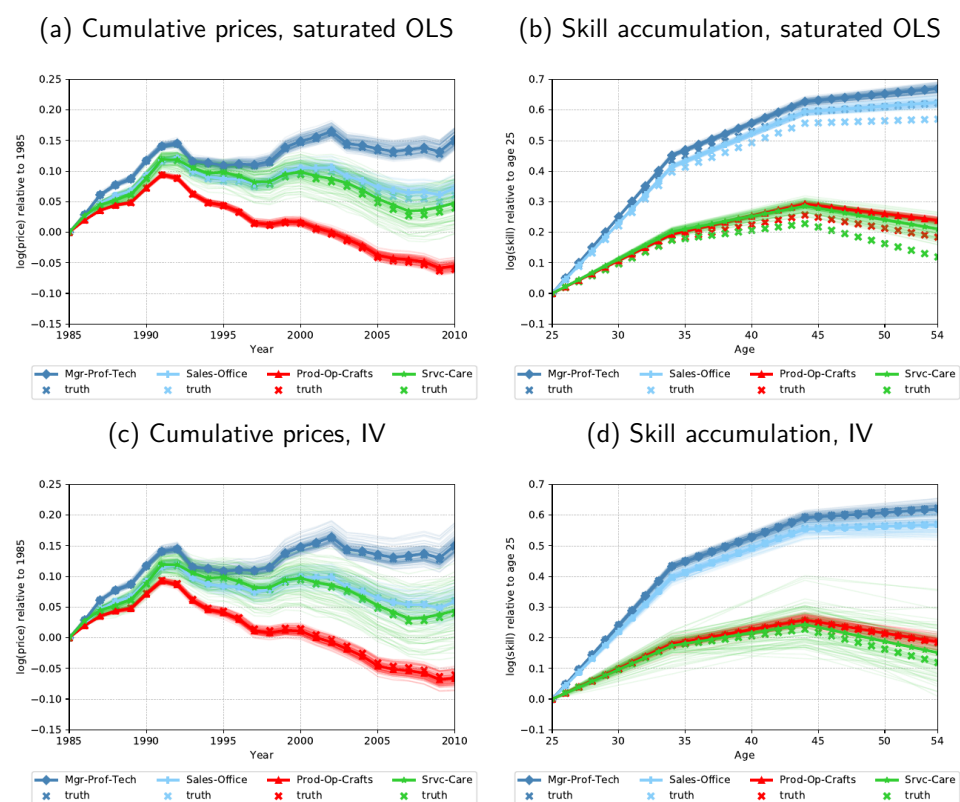


(d) Evolution of the wage distribution



Notes: The amount of stayers in the switching graphs corresponds to the area between the two dashed lines. The occupations are ordered by the average wage within the occupations across years 1985 until 2010. Entering or leaving an occupation to an occupation with a higher average wage is depicted by an area above the upper dashed line and vice versa for occupations with a lower average wage. The wage growth histogram was calculated using 100 equally sized bins between -0.5 and 0.5. The results are averages across the 100 Monte Carlo replications.

Figure 1.28: Estimation results, moderate switching costs and moderate shocks



Notes: Crosses “x” represent true values used to simulate workers’ wages. Solid lines are averages across the 100 Monte Carlo replications. Shimmering lines in the background represent individual Monte Carlo replications. OLS estimates as described by Equation (1.9). IV estimates as described at the end of Section 1.3.3.

Table 1.15: True and estimated skill accumulation parameters, saturated OLS

Previous sector	Current sector		Age group					
			[25, 34]		[35, 44]		[45, 54]	
			$\hat{\gamma}_{k',k,a}$	$\gamma_{k',k,a}^{true}$	$\hat{\gamma}_{k',k,a}$	$\gamma_{k',k,a}^{true}$	$\hat{\gamma}_{k',k,a}$	$\gamma_{k',k,a}^{true}$
Mgr-Prof-Tech	Mgr-Prof-Tech	γ	0.050	0.048	0.018	0.016	0.004	0.003
		σ_γ	0.000		0.000		0.000	
	Sales-Office	γ	0.126	0.063	0.045	0.009	-0.021	-0.010
		σ_γ	0.018		0.022		0.036	
	Prod-Op-Crafts	γ	0.027	0.023	-0.047	-0.011	-0.069	-0.022
		σ_γ	0.031		0.028		0.034	
	Srvc-Care	γ	-0.121	-0.008	-0.160	-0.036	-0.007	-0.004
		σ_γ	0.095		0.107		0.063	
Sales-Office	Mgr-Prof-Tech	γ	0.176	0.088	0.105	0.027	0.078	0.009
		σ_γ	0.009		0.014		0.019	
	Sales-Office	γ	0.046	0.044	0.018	0.016	0.003	0.001
		σ_γ	0.000		0.000		0.000	
	Prod-Op-Crafts	γ	0.122	0.056	0.087	0.019	0.029	-0.008
		σ_γ	0.018		0.016		0.025	
	Srvc-Care	γ	-0.041	0.010	-0.146	-0.034	-0.094	-0.024
		σ_γ	0.072		0.105		0.112	
Prod-Op-Crafts	Mgr-Prof-Tech	γ	0.173	0.075	0.127	0.042	0.103	0.021
		σ_γ	0.004		0.006		0.007	
	Sales-Office	γ	0.110	0.036	0.093	0.022	0.062	0.000
		σ_γ	0.010		0.010		0.013	
	Prod-Op-Crafts	γ	0.022	0.020	0.010	0.008	-0.005	-0.007
		σ_γ	0.000		0.000		0.000	
	Srvc-Care	γ	-0.038	-0.017	-0.056	-0.014	0.008	-0.009
		σ_γ	0.035		0.040		0.027	
Srvc-Care	Mgr-Prof-Tech	γ	0.194	0.099	0.159	0.063	0.134	0.041
		σ_γ	0.012		0.015		0.018	
	Sales-Office	γ	0.181	0.090	0.141	0.048	0.100	0.015
		σ_γ	0.014		0.020		0.034	
	Prod-Op-Crafts	γ	0.201	0.106	0.174	0.075	0.136	0.037
		σ_γ	0.010		0.011		0.014	
	Srvc-Care	γ	0.022	0.019	0.008	0.005	-0.008	-0.011
		σ_γ	0.001		0.001		0.001	

Notes: The table shows the estimated $\hat{\gamma}_{k',k,a}$, which is a scalar element of $\Gamma_{k',k}$ representing skill accumulation of age group a . The four groups are based on an aggregation of detailed occupations in the SIAB SUF as described in Appendix Table 1.4. OLS estimates as described by Equation (1.9).

Table 1.16: True and estimated skill accumulation parameters, IV

Previous sector	Current sector		Age group						
			[25, 34]		[35, 44]		[45, 54]		
			$\hat{\gamma}_{k',k,a}$	$\gamma_{k',k,a}^{true}$	$\hat{\gamma}_{k',k,a}$	$\gamma_{k',k,a}^{true}$	$\hat{\gamma}_{k',k,a}$	$\gamma_{k',k,a}^{true}$	
Mgr-Prof-Tech	Mgr-Prof-Tech	γ	0.048	0.048	0.016	0.016	0.003	0.003	
		σ_γ	0.001		0.001		0.001		
	Sales-Office	γ	0.377	0.063	0.393	0.009	0.369	-0.010	
		σ_γ	0.102		0.147		0.223		
	Prod-Op-Crafts	γ	0.496	0.023	0.432	-0.011	0.422	-0.022	
		σ_γ	0.123		0.111		0.129		
	Srvc-Care	γ	0.488	-0.008	0.506	-0.036	0.454	-0.004	
		σ_γ	0.528		0.496		0.333		
	Sales-Office	Mgr-Prof-Tech	γ	0.301	0.088	0.324	0.027	0.327	0.009
			σ_γ	0.062		0.131		0.250	
Sales-Office		γ	0.044	0.044	0.016	0.016	0.001	0.001	
		σ_γ	0.001		0.001		0.001		
Prod-Op-Crafts		γ	0.386	0.056	0.387	0.019	0.383	-0.008	
		σ_γ	0.095		0.111		0.163		
Srvc-Care		γ	0.513	0.010	0.476	-0.034	0.100	-0.024	
		σ_γ	0.472		0.530		0.838		
Prod-Op-Crafts		Mgr-Prof-Tech	γ	0.247	0.075	0.288	0.042	0.274	0.021
			σ_γ	0.065		0.090		0.126	
	Sales-Office	γ	0.314	0.036	0.316	0.022	0.311	0.000	
		σ_γ	0.131		0.098		0.163		
	Prod-Op-Crafts	γ	0.020	0.020	0.008	0.008	-0.007	-0.007	
		σ_γ	0.001		0.001		0.001		
	Srvc-Care	γ	0.416	-0.017	0.434	-0.014	0.365	-0.009	
		σ_γ	0.285		0.236		0.188		
	Srvc-Care	Mgr-Prof-Tech	γ	0.285	0.099	0.214	0.063	0.201	0.041
			σ_γ	0.132		0.171		0.139	
Sales-Office		γ	0.273	0.090	0.155	0.048	0.042	0.015	
		σ_γ	0.104		0.248		0.424		
Prod-Op-Crafts		γ	0.279	0.106	0.235	0.075	0.224	0.037	
		σ_γ	0.037		0.059		0.089		
Srvc-Care		γ	0.019	0.019	0.007	0.005	-0.009	-0.011	
		σ_γ	0.003		0.004		0.004		

Notes: The table shows the estimated $\hat{\gamma}_{k',k,a}$, which is a scalar element of $\Gamma_{k',k}$ representing skill accumulation of age group a . The four groups are based on an aggregation of detailed occupations in the SIAB SUF as described in Appendix Table 1.4. IV estimates as described at the end of Section 1.3.3.

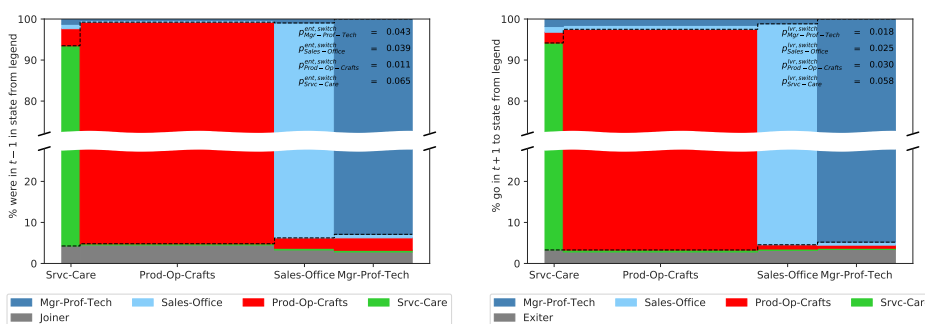
1.C.4.3 High Switching Costs, Highly Dispersed Shocks

Table 1.17: Parameters


Parameter	Value
N	50000
Repetitions	100
Skill shocks in k	observed wage growth distribution, $(\mu, \sigma_k) = (0, 1.5 \cdot \sigma_{\Delta \log(w_i)}^{SIAB})$
Stayers accumulation $\gamma_{k,k,a}, k' = k$	$\hat{\gamma}_{k,k,a}^{SIAB}$
Cross accumulation $\gamma_{k',k,a}, k' \neq k$	$\frac{1}{3} \hat{\gamma}_{k',k,a}^{SIAB}$
ρ in $\varepsilon_{i,t} = \rho \varepsilon_{i,t-1} + v_{i,t}$	0
Switching costs c	0.2
Amenity trends, $t = 1985, \dots, 2010$	$[\Delta \Psi_{k,t}]_{k=1,\dots,4} = [0, 0, 0, 0]$

Figure 1.29: Descriptives, high switching costs and highly dispersed shocks


(a) Occupation entrants/incumb. in $t + 1$ (b) Occupation leavers/stayers in $t - 1$



(c) Distribution of annual wage growth

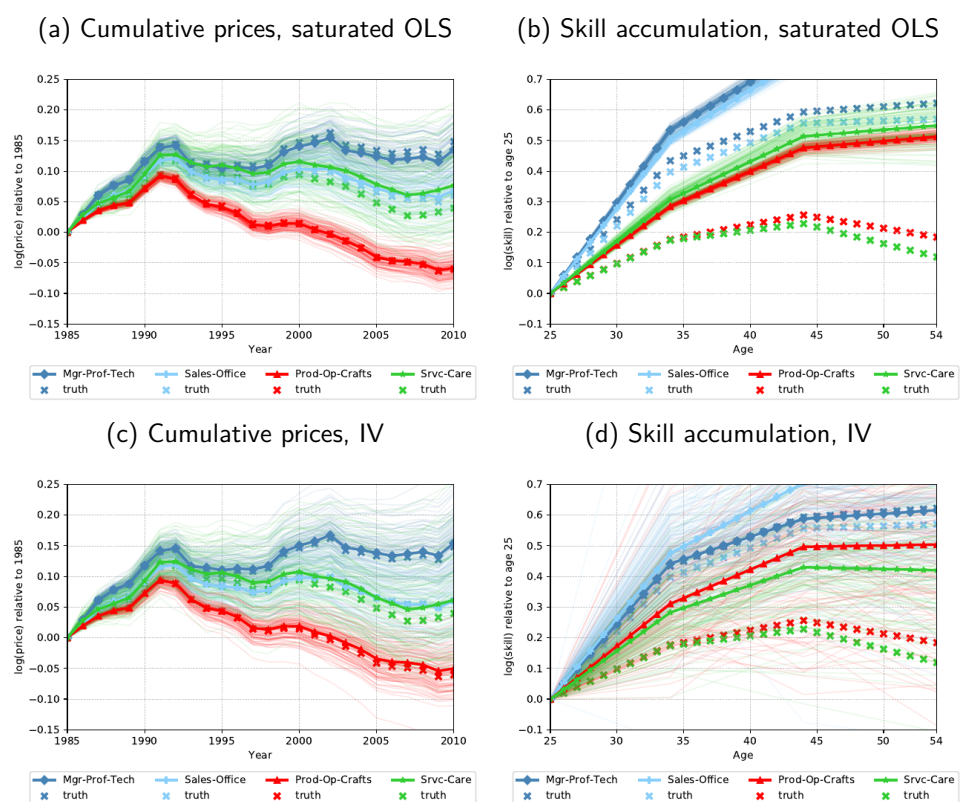


(d) Evolution of the wage distribution



Notes: The amount of stayers in the switching graphs corresponds to the area between the two dashed lines. The occupations are ordered by the average wage within the occupations across years 1985 until 2010. Entering or leaving an occupation to an occupation with a higher average wage is depicted by an area above the upper dashed line and vice versa for occupations with a lower average wage. The wage growth histogram was calculated using 100 equally sized bins between -0.5 and 0.5. The results are averages across the 100 Monte Carlo replications.

Figure 1.30: Estimation results, high switching costs and highly dispersed shocks



Notes: Crosses “x” represent true values used to simulate workers’ wages. Solid lines are averages across the 100 Monte Carlo replications. Shimmering lines in the background represent individual Monte Carlo replications. OLS estimates as described by Equation (1.9). IV estimates as described at the end of Section 1.3.3.

1.C.5 Changing Amenities

Here, we introduce trends in non-wage amenities to the worker's decision problem. These trends make some occupations more attractive over time even when prices did not change. We implement them so that relative to Prod-Op-Crafts, the other occupations became less attractive. This makes workers move into Prod-Op-Crafts despite falling prices. Figure 1.31 shows the descriptives.



Figure 1.32 shows the estimation results. The baseline method is biased as skill price and amenity values are confounded by each other, making us overpredict the fall in Srvc-Care prices where amenities fell. The adjustment described in Equation (1.28) takes care of this.

Table 1.18: Parameters


Parameter	Value
N	50000
Repetitions	100
Skill shocks in k	observed wage growth distribution, $(\mu, \sigma_k) = (0, 0.5 \cdot \sigma_{\Delta \log(w_i)}^{SIAB})$
Stayers accumulation $\gamma_{k,k,a}, k' = k$	$\hat{\gamma}_{k,k,a}^{SIAB}$
Cross accumulation $\gamma_{k',k,a}, k' \neq k$	$\frac{1}{3} \hat{\gamma}_{k',k,a}^{SIAB}$
ρ in $\varepsilon_{i,t} = \rho \varepsilon_{i,t-1} + v_{i,t}$	0
Switching costs c	0
Amenity trends, $t = 1985, \dots, 2010$	$[\Delta \Psi_{k,t}]_{k=1,\dots,4} = [0.02, 0, 0, 0]$

Figure 1.31: Descriptives, trends in amenities


(a) Occupation entrants/incumb. in $t + 1$ (b) Occupation leavers/stayers in $t - 1$



(c) Distribution of annual wage growth

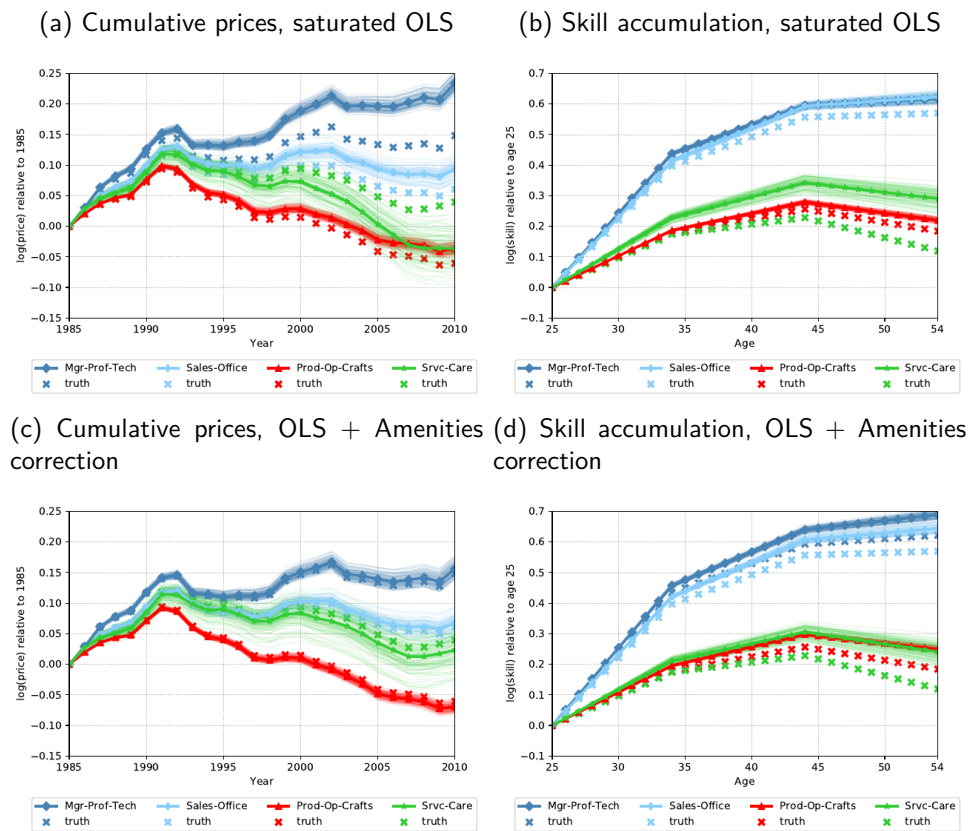


(d) Evolution of the wage distribution



Notes: The amount of stayers in the switching graphs corresponds to the area between the two dashed lines. The occupations are ordered by the average wage within the occupations across years 1985 until 2010. Entering or leaving an occupation to an occupation with a higher average wage is depicted by an area above the upper dashed line and vice versa for occupations with a lower average wage. The wage growth histogram was calculated using 100 equally sized bins between -0.5 and 0.5. The results are averages across the 100 Monte Carlo replications.

Figure 1.32: Estimation results, trends in amenities



Notes: Crosses “x” represent true values used to simulate workers’ wages. Solid lines are averages across the 100 Monte Carlo replications. Shimmering lines in the background represent individual Monte Carlo replications. OLS estimates as described by Equation (1.9). Estimates from the extended model which allows to control for changing amenity values of occupations as described by Equation (1.28). Amenities are allowed to vary by age group and identified relative to Prod-Op-Crafts.

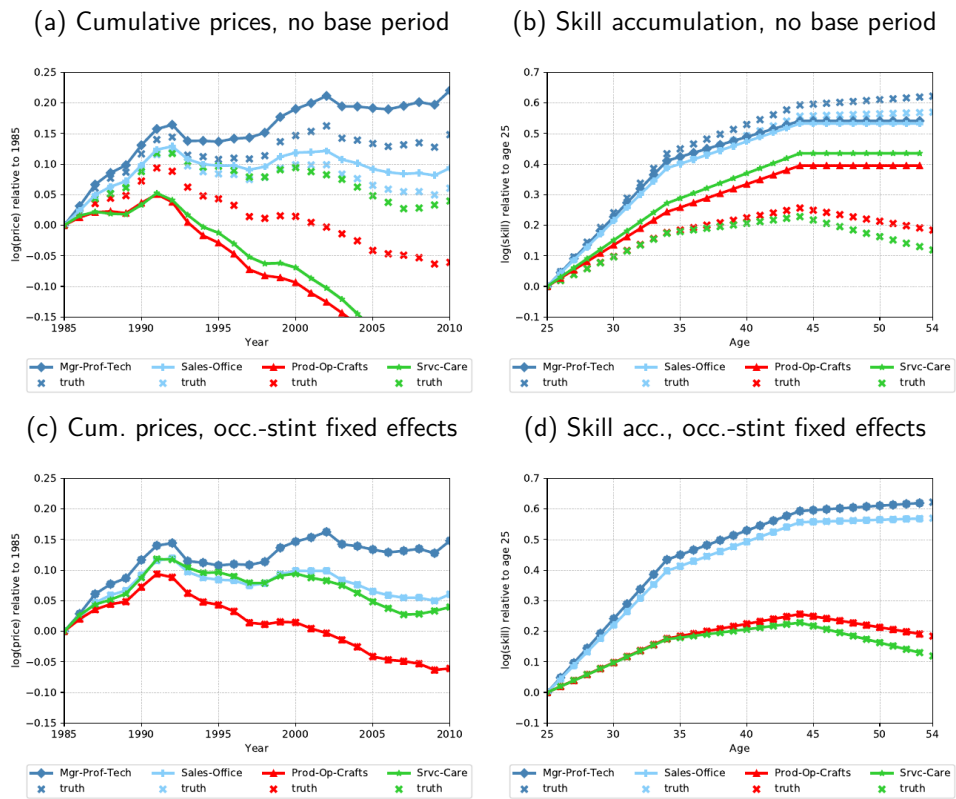
1.C.6 Occupation-Specific Fixed Effects

Finally, we compare our estimation method to an alternative approach proposed by Cortes (2016) who uses occupation-specific fixed effects to estimate changing skill prices. The top row of Figure 1.33 implements this approach without a base period. As discussed in Section 1.B.4, one age group within $\hat{\Gamma}_k$ is omitted because of perfect multicollinearity. Therefore, all the estimated parameters have to be interpreted relative to that age group's skill accumulation. This is quite complicated since some skill price changes are loaded on the estimated skill accumulation parameters and vice versa. Indeed, the top row of Figure 1.34 shows that both $\Delta\hat{\pi}_{k,t}$ and $\hat{\Gamma}_k$ substantially deviate from the truth and in the respective opposite direction for each occupation.

The bottom row of Figure 1.33 implements individuals' occupation-specific fixed effects in the way that we recommend, i.e., with a base period, occupation-*stint* specific fixed effects, and occupation-specific age profiles. This is very similar to the occupation-specific tenure profiles that Cortes (2016) uses in one of his key robustness checks (see again our discussion in Section 1.B.4) plus the base period. Supporting this specification, the skill prices and skill accumulation are perfectly identified if there are not any idiosyncratic skill shocks and therefore exogenous mobility holds.

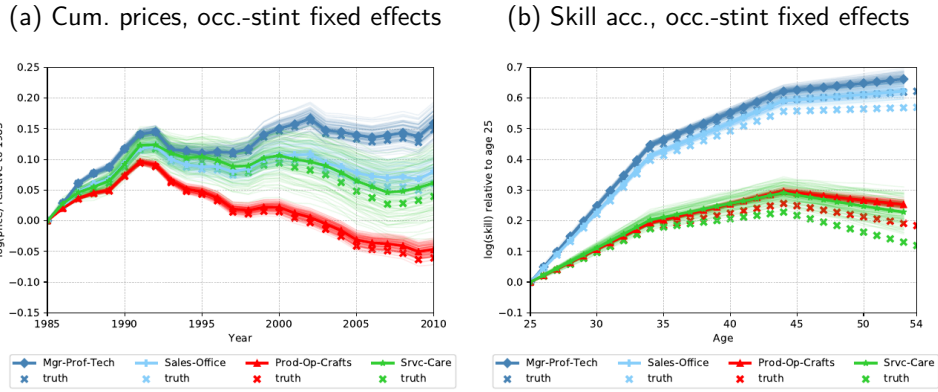
Next we add moderate skill shocks to the data generating process as detailed in Table 1.6 above. We run our recommended regression with occupation-stint specific fixed effects for this sample. Figure 1.34 depicts the results, showing that fixed effects approach still performs well. Finally, Figure 1.35 shows the estimation results for the data generating process with large skill shocks from Table 1.9. The prices are now substantially off for three out of four occupations, and the skill accumulation estimates are far away from the truth. This was predicted by us in the main text and in Section 1.B.4, as now endogenous occupation switching and *staying* becomes quantitatively important. It contrasts especially with the IV implementation of our approach in Figure 1.22, which comes very close to the true skill prices and reasonably close to skill accumulation even with large shocks.

Figure 1.33: Estimation results, no shocks as in Table 1.5



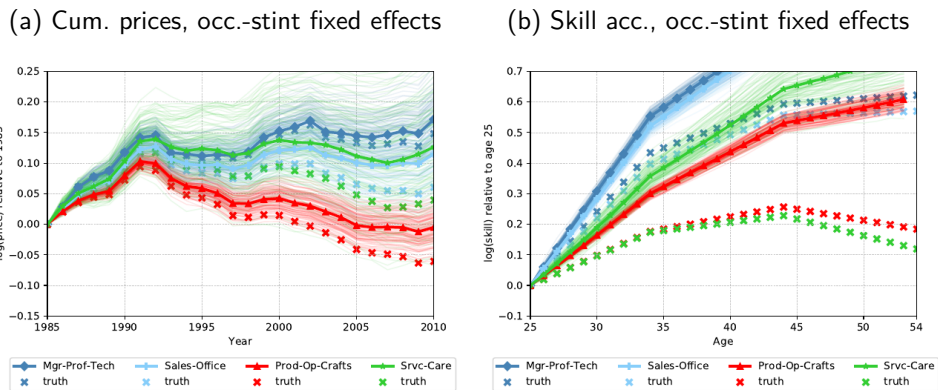
Notes: Crosses “x” represent true values used to simulate workers’ wages. Solid lines are averages across the 100 Monte Carlo replications. Shimmering lines in the background represent individual Monte Carlo replications. There is no variation in the background to be seen as there is no randomness included. Panels 1.33a and 1.33b show results from fixed effects estimation when no base period is included as described in Section 1.B.4.1. The remaining two Panels show estimates which are identified from year specific occupation fixed effects while including a separate worker-occupation fixed effects for each time the worker revisits an occupation (after a possible break or after return from another occupation). Additionally, we include controls for age and occupation dependent skill accumulation following Equation (1.39).

Figure 1.34: Estimation results, moderate shocks as in Table 1.6



Notes: Crosses “x” represent true values used to simulate workers’ wages. Solid lines are averages across the 100 Monte Carlo replications. Shimmering lines in the background represent individual Monte Carlo replications. Estimates are identified from year specific occupation fixed effects while including a separate worker-occupation fixed effects for each time the worker revisits an occupation (after a possible break or after return from another occupation). Additionally, we include controls for age and occupation dependent skill accumulation following Equation (1.39).

Figure 1.35: Estimation results, highly dispersed shocks as in Table 1.9



Notes: Crosses “x” represent true values used to simulate workers’ wages. Solid lines are averages across the 100 Monte Carlo replications. Shimmering lines in the background represent individual Monte Carlo replications. Estimates are identified from year specific occupation fixed effects while including a separate worker-occupation fixed effects for each time the worker revisits an occupation (after a possible break or after return from another occupation). Additionally, we include controls for age and occupation dependent skill accumulation following Equation (1.39).

1.D Additional Results for Section 1.4.1

1.D.1 Further Results of OLS Estimation

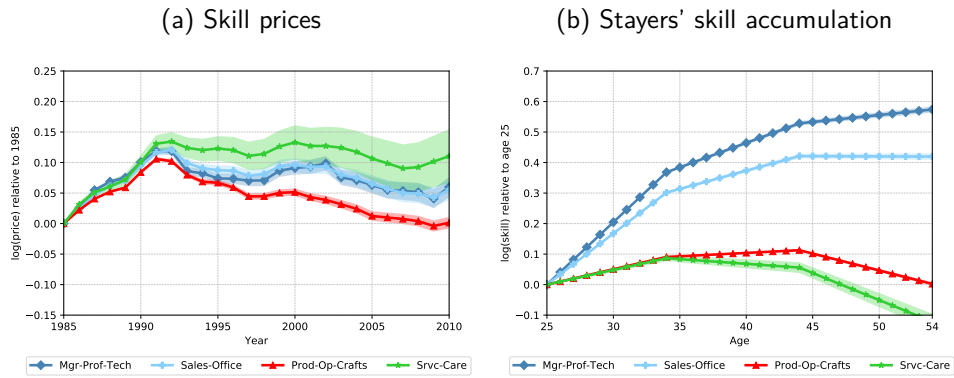
Table 1.19: Estimated skill accumulation coefficients (occupation groups, OLS)

Previous sector	Current sector		Age group			
			[25, 34]	[35, 44]	[45, 54]	
Mgr-Prof-Tech	Mgr-Prof-Tech	γ	0.048	0.016	0.003	
		σ_γ	0.000	0.000	0.000	
	Sales-Office	γ	0.141	0.012	-0.034	
		σ_γ	0.004	0.004	0.005	
	Prod-Op-Crafts	γ	0.020	-0.048	-0.068	
		σ_γ	0.005	0.005	0.006	
	Srvc-Care	γ	-0.071	-0.124	-0.016	
		σ_γ	0.011	0.012	0.016	
	Sales-Office	Mgr-Prof-Tech	γ	0.221	0.065	0.028
			σ_γ	0.003	0.004	0.005
Sales-Office		γ	0.044	0.016	0.001	
		σ_γ	0.000	0.000	0.000	
Prod-Op-Crafts		γ	0.125	0.042	-0.024	
		σ_γ	0.004	0.005	0.007	
Srvc-Care		γ	-0.014	-0.117	-0.074	
		σ_γ	0.010	0.012	0.015	
Prod-Op-Crafts		Mgr-Prof-Tech	γ	0.206	0.117	0.073
			σ_γ	0.003	0.003	0.005
	Sales-Office	γ	0.088	0.059	0.008	
		σ_γ	0.003	0.004	0.006	
	Prod-Op-Crafts	γ	0.020	0.008	-0.007	
		σ_γ	0.000	0.000	0.000	
	Srvc-Care	γ	-0.071	-0.050	-0.019	
		σ_γ	0.004	0.005	0.006	
	Srvc-Care	Mgr-Prof-Tech	γ	0.279	0.184	0.133
			σ_γ	0.009	0.011	0.016
Sales-Office		γ	0.250	0.139	0.055	
		σ_γ	0.008	0.011	0.015	
Prod-Op-Crafts		γ	0.300	0.220	0.123	
		σ_γ	0.004	0.006	0.007	
Srvc-Care		γ	0.019	0.005	-0.011	
		σ_γ	0.001	0.001	0.001	

Notes: The table shows the estimated $\hat{\gamma}_{k',k,a}$, which is a scalar element of $\Gamma_{k',k}$, representing skill accumulation of age group a . The four groups are based on an aggregation of detailed occupations in the SIAB SUF as described in Appendix Table 1.4. OLS estimates as described by Equation (1.9).

1.D.2 Instrumental Variable Estimates

Figure 1.36: The evolution of skill prices and skill accumulation of stayers



Notes: Panel 1.36a shows changes in skill price IV estimates over time as detailed in Section 1.3.3. Panel 1.36b shows stayers' skill accumulation profiles estimated with IV. Shaded lines in the background represent the 120 detailed occupations in the SIAB SUF. The four groups show an aggregation of these detailed occupations as described in Appendix Table 1.4. The thickness of a shaded background line corresponds to the number of employed workers in an occupation averaged across years 1985 until 2010. The shaded areas around the four lines are 95% confidence intervals.

Table 1.20: Estimated skill accumulation coefficients (occupation groups, IV)

Previous sector	Current sector		Age group		
			[25, 34]	[35, 44]	[45, 54]
Mgr-Prof-Tech	Mgr-Prof-Tech	γ	0.041	0.016	0.005
		σ_γ	0.001	0.000	0.001
	Sales-Office	γ	0.216	0.194	0.227
		σ_γ	0.058	0.055	0.081
	Prod-Op-Crafts	γ	-0.141	0.056	-0.116
		σ_γ	0.045	0.045	0.065
Srvc-Care	γ	-0.400	-0.062	-0.051	
	σ_γ	0.179	0.138	0.229	
Sales-Office	Mgr-Prof-Tech	γ	0.397	0.247	0.156
		σ_γ	0.055	0.040	0.061
	Sales-Office	γ	0.034	0.012	-0.000
		σ_γ	0.001	0.001	0.001
	Prod-Op-Crafts	γ	0.074	0.101	-0.119
		σ_γ	0.036	0.037	0.075
Srvc-Care	γ	-0.061	0.067	-0.284	
	σ_γ	0.168	0.185	0.291	
Prod-Op-Crafts	Mgr-Prof-Tech	γ	0.371	0.327	0.117
		σ_γ	0.054	0.055	0.092
	Sales-Office	γ	0.343	0.337	0.361
		σ_γ	0.071	0.070	0.109
	Prod-Op-Crafts	γ	0.010	0.002	-0.011
		σ_γ	0.000	0.000	0.000
Srvc-Care	γ	0.372	0.193	-0.114	
	σ_γ	0.095	0.107	0.176	
Srvc-Care	Mgr-Prof-Tech	γ	0.390	0.222	0.327
		σ_γ	0.192	0.129	0.109
	Sales-Office	γ	0.498	0.619	0.419
		σ_γ	0.185	0.240	0.243
	Prod-Op-Crafts	γ	0.198	0.346	0.216
		σ_γ	0.050	0.051	0.077
Srvc-Care	γ	0.010	-0.003	-0.018	
	σ_γ	0.002	0.001	0.001	

Notes: The table shows the estimated $\hat{\gamma}_{k',k,a}$, which is a scalar element of $\Gamma_{k',k}$, representing skill accumulation of age group a . The four groups are based on an aggregation of detailed occupations in the SIAB SUF as described in Appendix Table 1.4. IV estimates as described at the end of Section 1.3.3.

1.E Detailed Analysis of Skill Selection

1.E.1 Derivation of Equation (1.13)

$$\begin{aligned}
& \mathbb{E}[s_{k,i,t} | I_{k,i,t} = 1] - \mathbb{E}[s_{k,i,t-1} | I_{k,i,t-1} = 1] = \\
& \underbrace{\mathbb{E}[s_{k,i,t} | I_{k,i,t} = 1, I_{k,i,t-1} = 1]}_{\mathbb{E}[s_{k,i,t}^{incumb}]} \cdot \underbrace{P(I_{k,i,t-1} = 1 | I_{k,i,t} = 1)}_{1-p_{k,t}^{ent}} \\
& + \underbrace{\mathbb{E}[s_{k,i,t} | I_{k,i,t} = 1, I_{k,i,t-1} = 0]}_{\mathbb{E}[s_{k,i,t}^{ent}]} \cdot \underbrace{P(I_{k,i,t-1} = 0 | I_{k,i,t} = 1)}_{p_{k,t}^{ent}} \\
& - \left(\underbrace{\mathbb{E}[s_{k,i,t-1} | I_{k,i,t-1} = 1, I_{k,i,t} = 1]}_{\mathbb{E}[s_{k,i,t-1}^{sty}]} \cdot \underbrace{P(I_{k,i,t} = 1 | I_{k,i,t-1} = 1)}_{1-p_{k,t-1}^{lvr}} \right. \\
& \left. + \underbrace{\mathbb{E}[s_{k,i,t-1} | I_{k,i,t-1} = 1, I_{k,i,t} = 0]}_{\mathbb{E}[s_{k,i,t-1}^{lvr}]} \cdot \underbrace{P(I_{k,i,t} = 0 | I_{k,i,t-1} = 1)}_{p_{k,t-1}^{lvr}} \right)
\end{aligned} \tag{1.44}$$

First notice that period $t - 1$ stayers are the same individuals as period t incumbents and define $\mathbb{E}[\Delta s_{k,i,t}^{incumb}] \equiv \mathbb{E}[s_{k,i,t}^{incumb}] - \mathbb{E}[s_{k,i,t-1}^{sty}]$. We can now combine the second and fourth as well as the third and fifth row of (1.44):

$$\begin{aligned}
& \mathbb{E}[s_{k,i,t} | I_{k,i,t} = 1] - \mathbb{E}[s_{k,i,t-1} | I_{k,i,t-1} = 1] = \\
& (1 - p_{k,t-1}^{lvr}) \cdot \mathbb{E}[\Delta s_{k,i,t}^{incumb}] + (p_{k,t-1}^{lvr} - p_{k,t}^{ent}) \cdot \mathbb{E}[s_{k,i,t}^{incumb}] + \\
& p_{k,t-1}^{lvr} \cdot (\mathbb{E}[s_{k,i,t}^{ent}] - \mathbb{E}[s_{k,i,t-1}^{lvr}]) + (p_{k,t}^{ent} - p_{k,t-1}^{lvr}) \cdot \mathbb{E}[s_{k,i,t}^{ent}].
\end{aligned}$$

This decomposes the skill change with marginal selection only of entrants:

$$\begin{aligned}
& \mathbb{E}[s_{k,i,t} | I_{k,i,t} = 1] - \mathbb{E}[s_{k,i,t-1} | I_{k,i,t-1} = 1] = \\
& \underbrace{(1 - p_{k,t-1}^{lvr}) \cdot \mathbb{E}[\Delta s_{k,i,t}^{incumb}]}_{\text{1. Skill accumulation of } t-1 \text{ stayers}} \\
& + \underbrace{p_{k,t-1}^{lvr} \cdot (\mathbb{E}[s_{k,i,t}^{ent}] - \mathbb{E}[s_{k,i,t-1}^{lvr}])}_{\text{2. Churning of leavers: difference entrants in } t, \text{ leavers after } t-1} \\
& + \underbrace{(p_{k,t}^{ent} - p_{k,t-1}^{lvr}) \cdot (\mathbb{E}[s_{k,i,t}^{ent}] - \mathbb{E}[s_{k,i,t}^{incumb}])}_{\text{3. Marginal selection of entrants}}
\end{aligned} \tag{1.45}$$

The inverse way of factoring out

$$\begin{aligned}
& \mathbb{E}[s_{k,i,t} | I_{k,i,t} = 1] - \mathbb{E}[s_{k,i,t-1} | I_{k,i,t-1} = 1] = \\
& (1 - p_{k,t}^{ent}) \cdot \mathbb{E}[\Delta s_{k,i,t}^{incumb}] - (p_{k,t}^{ent} - p_{k,t-1}^{lvr}) \cdot \mathbb{E}[s_{k,i,t-1}^{sty}] + \\
& p_{k,t}^{ent} \cdot (\mathbb{E}[s_{k,i,t}^{ent}] - \mathbb{E}[s_{k,i,t-1}^{lvr}]) - (p_{k,t-1}^{lvr} - p_{k,t}^{ent}) \cdot \mathbb{E}[s_{k,i,t-1}^{lvr}],
\end{aligned}$$

yields the decomposition with marginal selection only of leavers:

$$\begin{aligned}
 & E[s_{k,i,t} | I_{k,i,t} = 1] - E[s_{k,i,t-1} | I_{k,i,t-1} = 1] = & (1.46) \\
 & \underbrace{(1 - p_{k,t-1}^{ent}) \cdot E[\Delta s_{k,i,t}^{incumb}]}_{\text{1. Skill accumulation of } t \text{ incumbents}} \\
 & + \underbrace{p_{k,t}^{ent} \cdot (E[s_{k,i,t}^{ent}] - E[s_{k,i,t-1}^{lvr}])}_{\text{2. Churning of entrants: difference entrants in } t, \text{ leavers after } t-1} \\
 & + \underbrace{(p_{k,t}^{ent} - p_{k,t-1}^{lvr}) \cdot (E[s_{k,i,t-1}^{lvr}] - E[s_{k,i,t-1}^{sty}])}_{\text{3. Marginal selection of leavers}}
 \end{aligned}$$

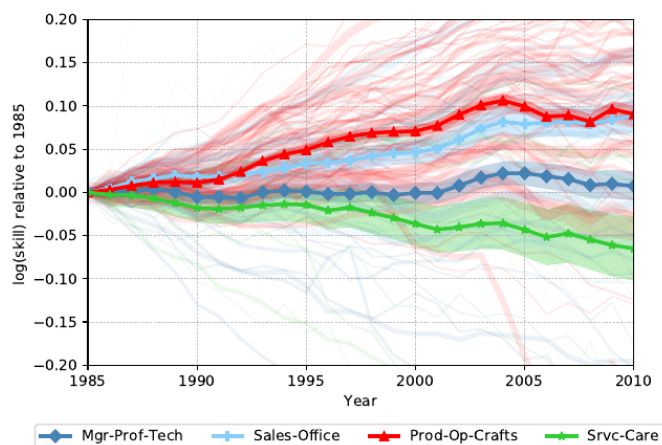
Adding (1.45) and (1.46) and dividing by 2 gives Equation (1.13).

Figure 1.8 shows this “average” decomposition with the mean of the respective 1. Skill accumulation, 2. Churning, and 3. Marginal selection effects. In Figure 1.39, we separately plot marginal selection of entrants (1.45, Panel a) and leavers (1.46, Panel b). The graphs reveal both to be of the same order of magnitude with marginal selection for entrants being somewhat steeper.

1.E.2 Skill Changes and the Marginal Selection Effect

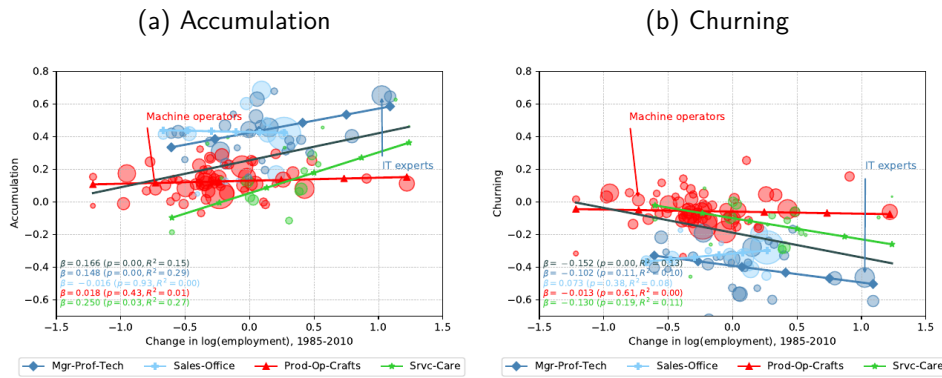
This section reports additional details for the skill changes in occupations. First, Figure 1.37 shows the skill trends implied by average wages minus the estimated skill prices following Equation (1.12).

Figure 1.37: The evolution of average skill by occupation



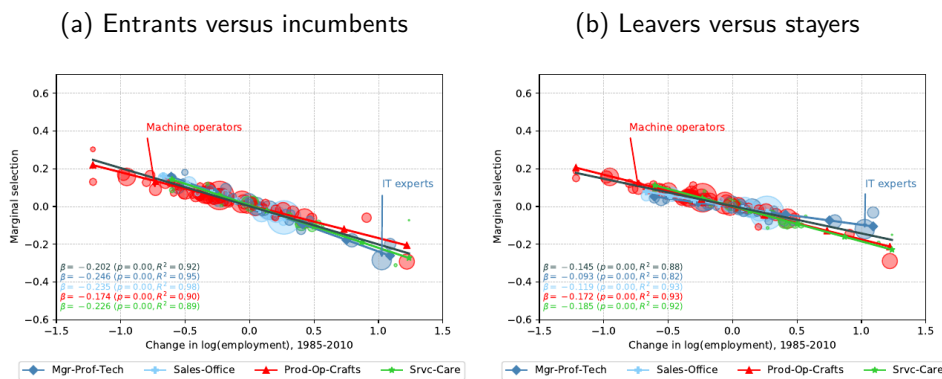
Notes: The figure shows estimated skill changes as implied in Equation (1.12). OLS estimates as described by Equation (1.9). Shaded lines in the background represent the 120 detailed occupations in the SIAB SUF. The four groups show an aggregation of these detailed occupations as described in Appendix Table 1.4. The thickness of a shaded background line corresponds to the number of employed workers in an occupation averaged across years 1985 until 2010. The shaded areas around the four lines are 95% confidence intervals.

Figure 1.38: Employment changes vs. accumulation and churning separately



Notes: Results correspond to sample averages following Equation (1.13). The horizontal axes in all panels show the change of the log number of employed workers within an occupation between 1985 and 2010. One bubble represents one of the 120 detailed occupations in the SIAB SUF. The four groups show an aggregation of these detailed occupations as described in Appendix Table 1.4. Bubble size corresponds to the number of employed workers in an occupation averaged across years 1985 until 2010. Regression lines across all occupations (black) and within the four broad groups (colored) are weighted by the number of employed workers.

Figure 1.39: Marginal selection of entrants and of leavers



Notes: Results from the separate decompositions (1.45) and (1.46). The horizontal axes in all panels show the change of the log number of employed workers within an occupation between 1985 and 2010. One bubble represents one of the 120 detailed occupations in the SIAB SUF. The four groups show an aggregation of these detailed occupations as described in Appendix Table 1.4. Bubble size corresponds to the number of employed workers in an occupation averaged across years 1985 until 2010. Regression lines across all occupations (black) and within the four broad groups (colored) are weighted by the number of employed workers.

1.E.3 Sources of the Marginal Selection Effect

This section investigates the sources of marginal selection in more detail. The effect can either be “classic selection” due to differences in skill endowments or it can be due to (differences in) skill changes of workers after they entered the occupation. These changes may be estimated skill accumulation but they will also have a deviating component which is consistent with workers who experience positive shocks endogenously staying in their occupation. We alternatively decompose the marginal selection effect into the contributions of occupation switchers, sample entrants and leavers, and movers from or into unemployment and out of the labor force.

The main text displays the formula for the average marginal selection effect in Equation (1.13). Section 1.E.1 in fact derived this by constructing marginal selection of entrants only $(p_{k,t}^{ent} - p_{k,t-1}^{lvr}) \cdot (E[s_{k,i,t}^{ent}] - E[s_{k,i,t}^{incumb}])$ in Equation (1.45) and of leavers only $(p_{k,t}^{ent} - p_{k,t-1}^{lvr}) \cdot (E[s_{k,i,t-1}^{lvr}] - E[s_{k,i,t-1}^{sty}])$ in Equation (1.46). This is what is separately plotted in Figure 1.39, and shown as separate contributions to the (average) marginal selection effect in Tables 1.1 and 1.2. In particular, the mean skills of incumbents are their mean log wages minus the estimated prices, i.e., $E[s_{k,i,t}^{incumb}] = E[w_{k,i,t}^{incumb}] - \pi_{k,t}$, and similarly for the skills of entrants $E[s_{k,i,t}^{ent}] = E[w_{k,i,t}^{ent}] - \pi_{k,t}$ and similarly for the skills of stayers $E[s_{k,i,t-1}^{sty}] = E[w_{k,i,t-1}^{sty}] - \pi_{k,t-1}$ and similarly for leavers $E[s_{k,i,t-1}^{lvr}] = E[w_{k,i,t-1}^{lvr}] - \pi_{k,t-1}$. Net entry $(p_{k,t}^{ent} - p_{k,t-1}^{lvr})$ is a constant that everything is multiplied with.

Table 1.1 breaks down the contributions to marginal selection by the origin or destination of marginal workers. In the top panel, we decompose the weighted mean skills of entrants into occupation group k :

$$E[s_{k,i,t}^{ent}] = p_{k,t}^{ent,swt} E[s_{k,i,t}^{ent,swt}] + p_{k,t}^{ent,unem} E[s_{k,i,t}^{ent,unem}] + p_{k,t}^{ent,olf} E[s_{k,i,t}^{ent,olf}] + p_{k,t}^{ent,smpl} E[s_{k,i,t}^{ent,smpl}], \quad (1.47)$$

where the shares of entrants who are occupation switchers $p_{k,t}^{ent,swt}$, entering from unemployment $p_{k,t}^{ent,unem}$ or out of the labor force during their careers $p_{k,t}^{ent,olf}$, and new sample entrants $p_{k,t}^{ent,smpl}$ sum to one. Accordingly, for leavers from k in the bottom panel of Table 1.1:

$$E[s_{k,i,t-1}^{lvr}] = p_{k,t-1}^{lvr,swt} E[s_{k,i,t-1}^{lvr,swt}] + p_{k,t-1}^{lvr,unem} E[s_{k,i,t-1}^{lvr,unem}] + p_{k,t-1}^{lvr,olf} E[s_{k,i,t-1}^{lvr,olf}] + p_{k,t-1}^{lvr,smpl} E[s_{k,i,t-1}^{lvr,smpl}] \quad (1.48)$$

In the case of direct occupation switchers, the means $E[s_{k,i,t}^{ent,swt}]$ and $E[s_{k,i,t-1}^{lvr,swt}]$ are further decomposed into entrants from and leavers to the four broad occupation groups.

Table 1.2 shows the contributions to marginal selection by the sources of workers' skills. We employ the longitudinal information in the data to separate workers' skill endowment at the most recent entry from their predicted skill accumulation and idiosyncratic deviations during the stay in the current occupation. In particular, we write the skills of a worker i in occupation k as:

$$s_{k,i,t} = s_{k,i,t_{k,i,0}} + \sum_{\tau=t_{k,i,0}}^{t-1} X'_{i,\tau} \hat{\Gamma}_{k,k} + \sum_{\tau=t_{k,i,0}+1}^t \hat{u}_{k,i,\tau}, \quad (1.49)$$

where the first term is the initial “endowment” when the worker entered this occupation at time $t_{k,i,0}$, the second term is predicted skill accumulation up to the current period

t , and the last term are the cumulated estimated skill shocks in k since entry for this particular worker.⁶⁰ Notice that, as discussed in Section 1.3.3, the $\hat{\Gamma}_{k,k,s}$ are estimated skill changes. They likely encompass structural accumulation parameters plus average skill shocks conditional on the choice of staying; especially in the OLS, which is used in Table 1.2. The estimated shocks $\hat{u}_{k,i,\tau}$ are therefore best interpreted simply as idiosyncratic deviations from the prediction $\hat{\Gamma}_{k,k}$.

The marginal selection for entering workers only can then be decomposed as

$$\begin{aligned}
 & \underbrace{(p_{k,t}^{ent} - p_{k,t-1}^{lvr}) \cdot (\mathbb{E}[s_{k,i,t}^{ent}] - \mathbb{E}[s_{k,i,t}^{incumb}])}_{\text{3. marginal selection for entrants}} = & (1.50) \\
 & \underbrace{(p_{k,t}^{ent} - p_{k,t-1}^{lvr}) (\mathbb{E}[s_{k,i,t}^{ent}] - \mathbb{E}[s_{k,i,t_{k,i,0}}^{incumb}])}_{\text{differences in endowments}} \\
 & + \underbrace{(p_{k,t}^{ent} - p_{k,t-1}^{lvr}) \left(\mathbb{E} \left[- \sum_{\tau=t_{k,i,0}}^{t-1} X'_{i,\tau} \Gamma_{k,k} | incumb \right] \right)}_{\text{systematic skill accumulation}} \\
 & + \underbrace{(p_{k,t}^{ent} - p_{k,t-1}^{lvr}) \left(\mathbb{E} \left[- \sum_{\tau=t_{k,i,0}+1}^t u_{k,i,\tau} | incumb \right] \right)}_{\text{idiosyncratic skill shocks}}
 \end{aligned}$$

⁶⁰Strictly speaking, we do not know levels of skill prices and skills but we can compute i 's overall accumulation $s_{k,i,t} - s_{k,i,t_{k,i,0}} = w_{k,i,t} - w_{k,i,t_{k,i,0}} - (\pi_{k,t} - \pi_{k,t_{k,i,0}})$ and use $\sum_{\tau=t_{k,i,0}}^{t-1} X'_{i,\tau} \hat{\Gamma}_{k,k}$ to back out $\sum_{\tau=t_{k,i,0}+1}^t \hat{u}_{k,i,\tau}$ from (1.49). Then, for comparisons of entrants versus incumbents or leavers versus stayers at a given point in time, levels of skill prices and thereby level shifters of skills in the population cancel out.

Notice however that the empirical implementation of (1.49) is not invariant to the more general acceleration/deceleration interpretation of the skill price ($\Delta \hat{\pi}_{k,t} = \Delta \pi_{k,t} - \Delta \pi_{k,base}$) and skill accumulation estimates ($\hat{\Gamma}_{k,k} = \Gamma_{k,k} + \Delta \pi_{k,base}$ for stayers). The reason is that our calculations then give us $\hat{s}_{k,i,t} - \hat{s}_{k,i,t_{k,i,0}} = s_{k,i,t} - s_{k,i,t_{k,i,0}} + (t - t_{k,i,0}) \Delta \pi_{k,base}$ and $\sum_{\tau=t_{k,i,0}}^{t-1} X'_{i,\tau} \hat{\Gamma}_{k,k} = \sum_{\tau=t_{k,i,0}}^{t-1} X'_{i,\tau} \Gamma_{k,k} + (t - t_{k,i,0}) \Delta \pi_{k,base}$. Because the resulting $(t - t_{k,i,0}) \Delta \pi_{k,base}$ on each side of (1.49) cancel out, the idiosyncratic deviations term remains nonetheless unaffected.

and correspondingly for leavers only

$$\begin{aligned}
 & \underbrace{(p_{k,t}^{ent} - p_{k,t-1}^{lvr}) \cdot \left(\mathbb{E}[s_{k,i,t-1}^{lvr}] - \mathbb{E}[s_{k,i,t-1}^{sty}] \right)}_{\text{3. marginal selection for leavers}} = & (1.51) \\
 & \underbrace{(p_{k,t}^{ent} - p_{k,t-1}^{lvr}) \left(\mathbb{E}[s_{k,i,t_{k,i},0}^{lvr}] - \mathbb{E}[s_{k,i,t_{k,i},0}^{sty}] \right)}_{\text{differences in endowments}} \\
 & + \underbrace{(p_{k,t}^{ent} - p_{k,t-1}^{lvr}) \left(\mathbb{E} \left[\sum_{\tau=t_{k,i},0}^{t-2} X'_{i,\tau} \Gamma_{k,k} |lvr \right] - \mathbb{E} \left[\sum_{\tau=t_{k,i},0}^{t-2} X'_{i,\tau} \Gamma_{k,k} |sty \right] \right)}_{\text{differences in systematic skill accumulation}} \\
 & + \underbrace{(p_{k,t}^{ent} - p_{k,t-1}^{lvr}) \left(\mathbb{E} \left[\sum_{\tau=t_{k,i},0+1}^{t-1} u_{k,i,\tau} |lvr \right] - \mathbb{E} \left[\sum_{\tau=t_{k,i},0+1}^{t-1} u_{k,i,\tau} |sty \right] \right)}_{\text{differences in idiosyncratic skill shocks}}.
 \end{aligned}$$

1.F Robustness of Estimated Price and Skill Changes

Section 1.4 of the main text has estimated the skill accumulation functions and changes in skill prices for detailed occupations as well as broader groups. We found that skill prices in fact increased with employment growth in Germany during 1985–2010, contrary to changes in average wages across occupations, and that marginal selection accounts for much of the systematic skill changes implied by the estimation. This section shows that these results are robust to various alternative sample definitions and estimation specifications.

1.F.1 Alternative Samples

A key robustness check is to allow for endogenous unemployment and exit from the labor force. In the main estimation we have assumed that coming into and exiting the sample is exogenous. This is obvious for individuals who reach age 25 or 54 (the borders of our sample age range) but it might not be an innocuous assumption during the career. In particular, workers may *choose to* become unemployed or exit the labor force if they obtain a sufficiently bad idiosyncratic skill shock or vice versa for a sufficiently good shock, and if the (time-limited) benefits or other non-labor income they obtain are sufficiently high. Our model would then be misspecified with unclear effects for the consistency of our estimates.

In Figure 1.40, we therefore assume that becoming unemployed or leaving the labor force temporarily is fully endogenous.⁶¹ We do this by imputing workers' wages and their occupation choices if they are unemployed or out of the labor force for any number of years between two spells of employment. We impute those by comparing pre and post non-employment wages and assigning workers the lower of those two wages adjusted for

⁶¹The reality is likely somewhere in between these two extremes. We do maintain the assumption that permanently leaving employment is exogenous because for prime age men this is quite rare (roughly 1.1% each year as opposed to 2.3% for temporary unemployment) and likely often due to relatively exogenous factors such as illness/death, moving to East Germany or abroad, becoming self-employed or civil servant, etc.

inflation. That is, we assume that workers could well have worked in the lower paying occupation but chose to become unemployed or exit the labor force for some period of time instead (further details in Section 1.A.1). On this sample, which is about 10% larger in size, we then repeat the estimation.

As mentioned in the main text already, the cross-accumulation parameters when workers switch with intermittent non-employment spells, especially into Prod-Op-Crafts and Srvc-Care occupations, tend to be lower in Table 1.21 than in Table 1.19 above. The correlation between wage and employment growth is approximately zero in Figure 1.40, but it is strongly positive between price and employment growth. A slightly flatter slope is induced by some fast growing occupations with many entrants from unemployment or out of the labor force spells whose estimated price growth diminishes when filling up those non-employment spells.⁶² Finally, the implied skill changes are again negative and quite closely related to marginal selection.

We have restricted our main sample to West German men as these can be defined consistently over the 1975–2010 period and many potentially confounding factors that may have affected women or foreigners, such as higher labor force participation, declining workplace discrimination (e.g., Hsieh et al., 2019), and rapidly rising educational attainment, do not apply. Nonetheless, the entry of women and foreigners as well as the reunification with the East constituted major supply shifts affecting the German labor market during our sample period. If women or foreigners were more inclined to work in Srvc-Care, for example, rising employment and falling wages in these occupations may be due to changes in labor supply. Also, if women or foreigners tend to earn less in certain occupations, estimated skill prices may be confounded by the closing of such gender or racial wage gaps over time. We therefore examine whether general equilibrium and composition effects due to supply shifts are important by checking if our estimates differ when we include these groups in our sample.

Figure 1.41 shows that skill prices hardly change when we include everyone, that is, women (increases sample by circa 69%), foreigners (6%), and individuals working in East Germany (15%), in the estimation. The implied skill changes and the marginal selection effect are somewhat steeper than in our main sample but qualitatively the same. Still more notable, when we estimate our model for prime aged West German women only, the relationship between occupations' skill price and employment growth is even stronger than for prime age men (i.e., 0.19 versus 0.15 slope of the regression line in Figure 1.42) and skill prices similarly tend to polarize (i.e., rise for the Mgr-Prof-Tech, Sales-Office, and Srvc-Care occupation groups). The same is true when dropping all workers whose nationality changes over the life cycle (Figure 1.43).

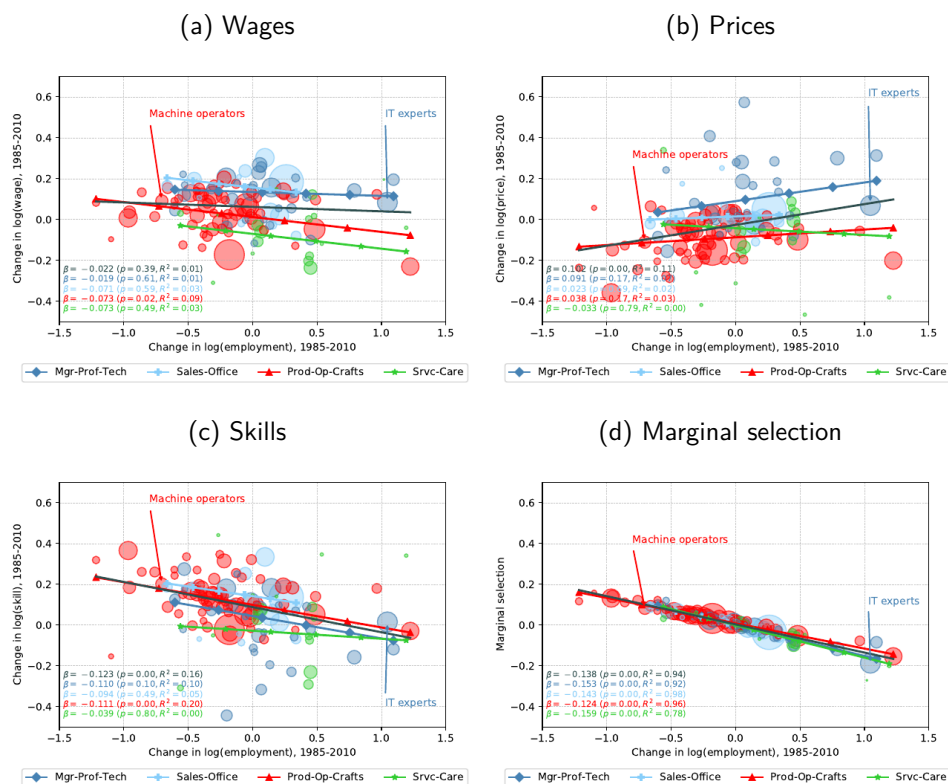
It is interesting to see that these results are similar despite a substantially different employment structure, with many more Sales-Office and Srvc-Care occupations among women than among men (becoming visible in the different bubble sizes). The results indicate that occupational demand shifts have largely driven the employment and skill price changes also for women, foreigners, and East Germans; apparently dominating other forces that may have worked on these demographic groups' changing labor market

⁶²One example are the “assistants” also discussed in the main text. They constitute a fairly low earning group with increasing turnover during the sample period. Instead of moving into that occupation, many workers might prefer to become unemployed or leave the labor force. Hence, we (increasingly) fill non-employment spells of later entrants to the assistants occupation up with low wages. This translates into lower price growth compared to the baseline sample. However, in total, these effects are not strong enough to substantially influence our estimates.

outcomes.

Finally, we widen the age range to 20–60 year old males. The results, depicted in Figure 1.44, are largely similar to our main prime age sample, with somewhat steeper slopes but also stronger marginal selection. The latter makes sense if labor market entrants in their early twenties were even less skilled compared to incumbents and lower-skilled workers were more likely to retire early.

Figure 1.40: Unemployment and leaving the labor force as a choice, i.e., filled non-employment spells



Notes: Unemployment and out of labor force spells are imputed by comparing the (real) wage after a non-employment spell with the wage before the non-employment spell. We then fill up the wage while in non-employment as the lower of those two wages adjusted for inflation and set the occupation within this time to the occupation that corresponds to that lower wage. See Appendix 1.A.1.4 for the details. The horizontal axes in all panels show the change of the log number of employed workers within an occupation between 1985 and 2010. One bubble represents one of the 120 detailed occupations in the SIAB SUF. The four groups show an aggregation of these detailed occupations as described in Appendix Table 1.4. Bubble size corresponds to the number of employed workers in an occupation averaged across years 1985 until 2010. Regression lines across all occupations (black) and within the four broad groups (colored) are weighted by the number of employed workers.

Table 1.21: Estimated skill accumulation coefficients (occupation groups, OLS), filled non-employment spells

Previous sector	Current sector		Age group			
			[25, 34]	[35, 44]	[45, 54]	
Mgr-Prof-Tech	Mgr-Prof-Tech	γ	0.045	0.015	0.002	
		σ_γ	0.000	0.000	0.000	
	Sales-Office	γ	0.093	-0.042	-0.135	
		σ_γ	0.004	0.004	0.005	
	Prod-Op-Crafts	γ	-0.043	-0.167	-0.210	
		σ_γ	0.005	0.005	0.006	
	Srvc-Care	γ	-0.238	-0.398	-0.401	
		σ_γ	0.011	0.011	0.015	
	Sales-Office	Mgr-Prof-Tech	γ	0.246	0.062	0.001
			σ_γ	0.003	0.004	0.005
Sales-Office		γ	0.041	0.014	0.000	
		σ_γ	0.000	0.000	0.000	
Prod-Op-Crafts		γ	0.030	-0.050	-0.156	
		σ_γ	0.004	0.005	0.006	
Srvc-Care		γ	-0.225	-0.331	-0.279	
		σ_γ	0.008	0.010	0.013	
Prod-Op-Crafts		Mgr-Prof-Tech	γ	0.250	0.161	0.066
			σ_γ	0.003	0.004	0.005
	Sales-Office	γ	0.047	0.067	-0.028	
		σ_γ	0.003	0.004	0.006	
	Prod-Op-Crafts	γ	0.016	0.007	-0.008	
		σ_γ	0.000	0.000	0.000	
	Srvc-Care	γ	-0.237	-0.185	-0.144	
		σ_γ	0.004	0.004	0.005	
	Srvc-Care	Mgr-Prof-Tech	γ	0.405	0.290	0.171
			σ_γ	0.009	0.011	0.016
Sales-Office		γ	0.211	0.162	0.049	
		σ_γ	0.008	0.010	0.015	
Prod-Op-Crafts		γ	0.259	0.186	0.106	
		σ_γ	0.004	0.005	0.007	
Srvc-Care		γ	0.013	0.004	-0.010	
		σ_γ	0.001	0.001	0.001	

Notes: The table shows the estimated $\hat{\gamma}_{k',k,a}$ for age groups a . k' is last period's occupation. k is the current occupation. Unemployment and out of labor force spells are imputed by comparing the (real) wage after a non-employment spell with the wage before the non-employment spell. We then fill up the wage while in non-employment as the lower of those two wages adjusted for inflation and set the occupation within this time to the occupation that corresponds to that lower wage. See Appendix 1.A.1.4 for the details.

Table 1.22: Contributions to marginal selection by origin and destination activities, filled non-employment spells

		Mgr- Prof- Tech	Sales- Office	Prod- Op- Crafts	Srvc- Care
Entrants	Switchers from Mgr-Prof-Tech		-0.06	0.01	0.00
	Switchers from Sales-Office	0.03		0.06	0.02
	Switchers from Prod-Op-Crafts	0.24	0.21		0.15
	Switchers from Srvc-Care	0.02	0.03	0.08	
	From unemployment	-0.00	-0.00	0.00	-0.00
	From outside of the labor force	0.05	0.02	0.09	0.00
	Sample entrants	0.43	0.55	0.49	0.43
Leavers	Switchers to Mgr-Prof-Tech		-0.02	0.02	0.03
	Switchers to Sales-Office	0.07		0.05	0.04
	Switchers to Prod-Op-Crafts	0.14	0.21		0.23
	Switchers to Srvc-Care	0.03	0.05	0.06	
	To unemployment	-0.00	-0.00	0.00	-0.00
	To outside of the labor force	0.24	0.21	0.36	0.26
	Sample leavers after age 54	-0.09	-0.07	-0.05	-0.03

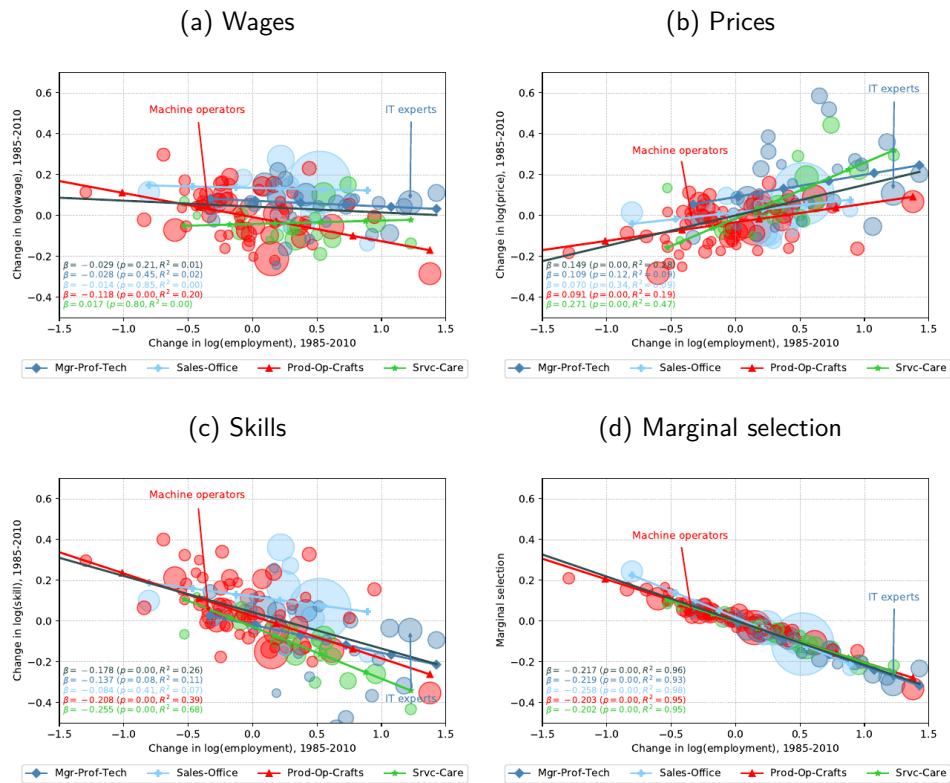
Notes: The numbers represent relative contributions to the marginal selection effect within each broad occupation group during 1985–2010. Columns sum to one. The explicit decomposition formulas are in Appendix 1.E.3.

Table 1.23: Contributions to marginal selection by source of skills, filled non-employment spells

		Mgr- Prof- Tech	Sales- Office	Prod- Op- Crafts	Srv- Care
Entrants	Endowment at the most recent entry into the occupation group	0.16	0.16	0.16	0.41
	Predicted skill accumulation since the most recent entry	0.41	0.46	0.28	0.07
	Deviation of skills from the prediction since the most recent entry	0.10	0.05	0.17	0.04
Leavers	Endowment at the most recent entry into the occupation group	0.00	-0.06	0.14	0.29
	Predicted skill accumulation since the most recent entry	0.18	0.25	0.03	0.10
	Deviation of skills from the prediction since the most recent entry	0.15	0.14	0.22	0.09

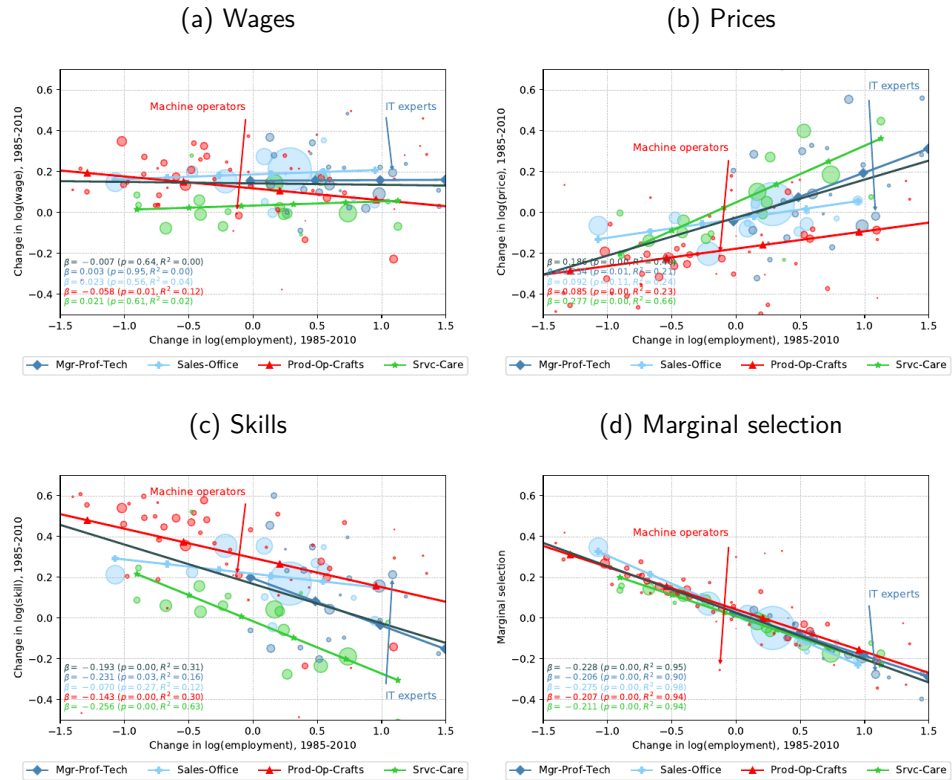
Notes: The numbers represent relative contributions to the marginal selection effect within each broad occupation group during 1985–2010. Columns sum to one. The explicit decomposition formulas are in Appendix 1.E.3.

Figure 1.41: Including East Germans, foreigners, and women



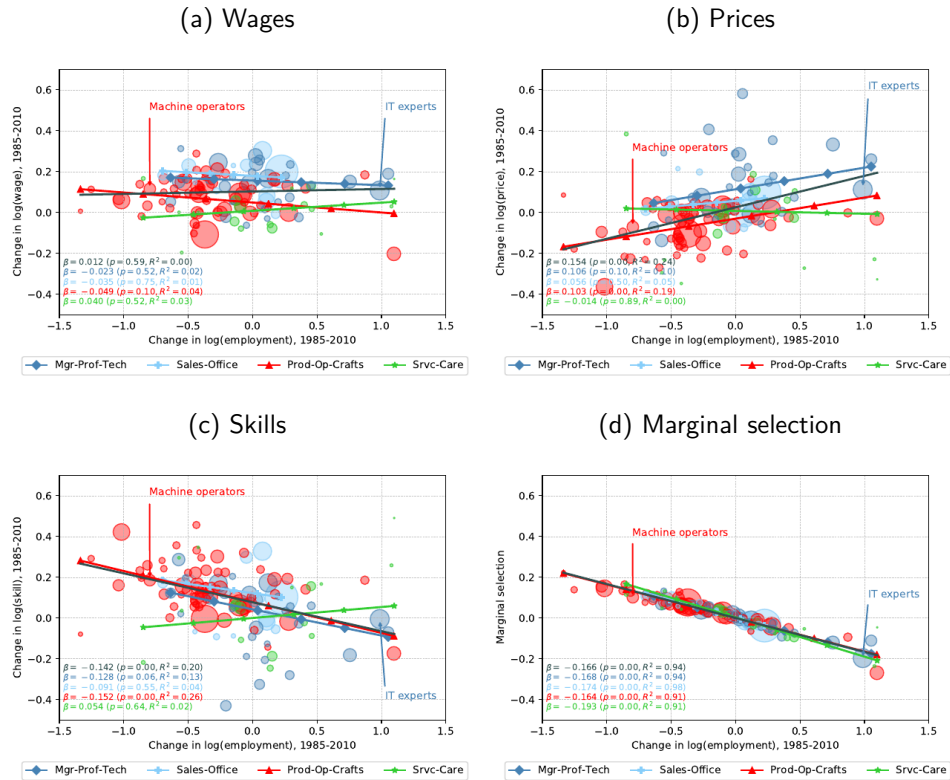
Notes: The sample additionally includes East Germans, foreigners and women. The horizontal axes in all panels show the change of the log number of employed workers within an occupation between 1985 and 2010. One bubble represents one of the 120 detailed occupations in the SIAB SUF. The four groups show an aggregation of these detailed occupations as described in Appendix Table 1.4. Bubble size corresponds to the number of employed workers in an occupation averaged across years 1985 until 2010. Regression lines across all occupations (black) and within the four broad groups (colored) are weighted by the number of employed workers.

Figure 1.42: Women only



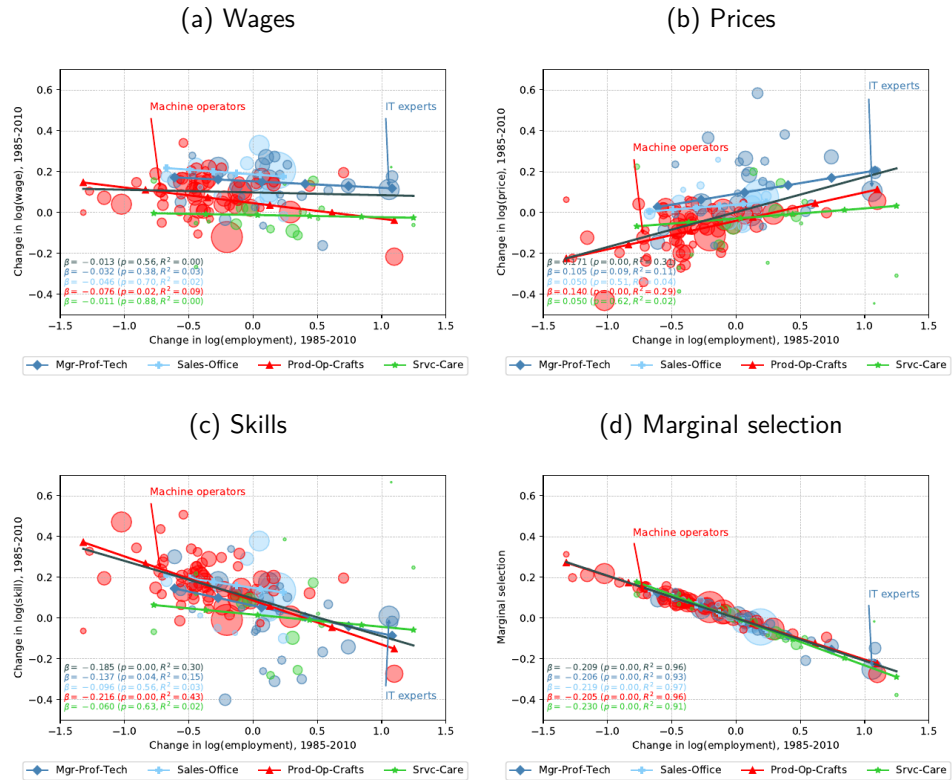
Notes: The sample is restricted to (full-time working) women. The horizontal axes in all panels show the change of the log number of employed workers within an occupation between 1985 and 2010. One bubble represents one of the 120 detailed occupations in the SIAB SUF. The four groups show an aggregation of these detailed occupations as described in Appendix Table 1.4. Bubble size corresponds to the number of employed workers in an occupation averaged across years 1985 until 2010. Regression lines across all occupations (black) and within the four broad groups (colored) are weighted by the number of employed workers.

Figure 1.43: Excluding anybody ever coded as a foreigner



Notes: The sample is the same as the baseline sample except that we also drop workers which are reported to be foreigners at some point in time. This includes, for instance, workers acquiring the German nationality at some later point in the life cycle. The horizontal axes in all panels show the change of the log number of employed workers within an occupation between 1985 and 2010. One bubble represents one of the 120 detailed occupations in the SIAB SUF. The four groups show an aggregation of these detailed occupations as described in Appendix Table 1.4. Bubble size corresponds to the number of employed workers in an occupation averaged across years 1985 until 2010. Regression lines across all occupations (black) and within the four broad groups (colored) are weighted by the number of employed workers.

Figure 1.44: All ages, 20–60 year olds



Notes: The sample is restricted to 20–60 year old men. The horizontal axes in all panels show the change of the log number of employed workers within an occupation between 1985 and 2010. One bubble represents one of the 120 detailed occupations in the SIAB SUF. The four groups show an aggregation of these detailed occupations as described in Appendix Table 1.4. Bubble size corresponds to the number of employed workers in an occupation averaged across years 1985 until 2010. Regression lines across all occupations (black) and within the four broad groups (colored) are weighted by the number of employed workers.

1.F.2 Alternative Estimation Specifications

This section presents the results from various extensions to our baseline method. Except for the flat-spot identification approach, all these extensions are applied to the main sample of prime aged, West German, full-time working men.

Noting that skill accumulation is rather flat for 45–54 year olds (e.g., Figure 1.6), we then apply Heckman, Lochner, et al. (1998)'s flat spot identification strategy to this subsample by setting skill accumulation to zero in all occupations. We see in Panel 1.45a that wage growth is again uncorrelated with employment growth in this older subsample whereas in Panel 1.45b skill price growth once more increases with occupations' size growth. This positive relationship is even stronger than in the full sample (regression slope of 0.18 compared to 0.15). Marginal selection does not fully account for the systematic skill changes anymore but still for more than half (also, the assumption of zero skill changes might not be strictly correct). Nonetheless, these results indicate that also in samples where dynamic considerations should not be a large concern, and when skill accumulation is arguably more or less constant, we get similar results.

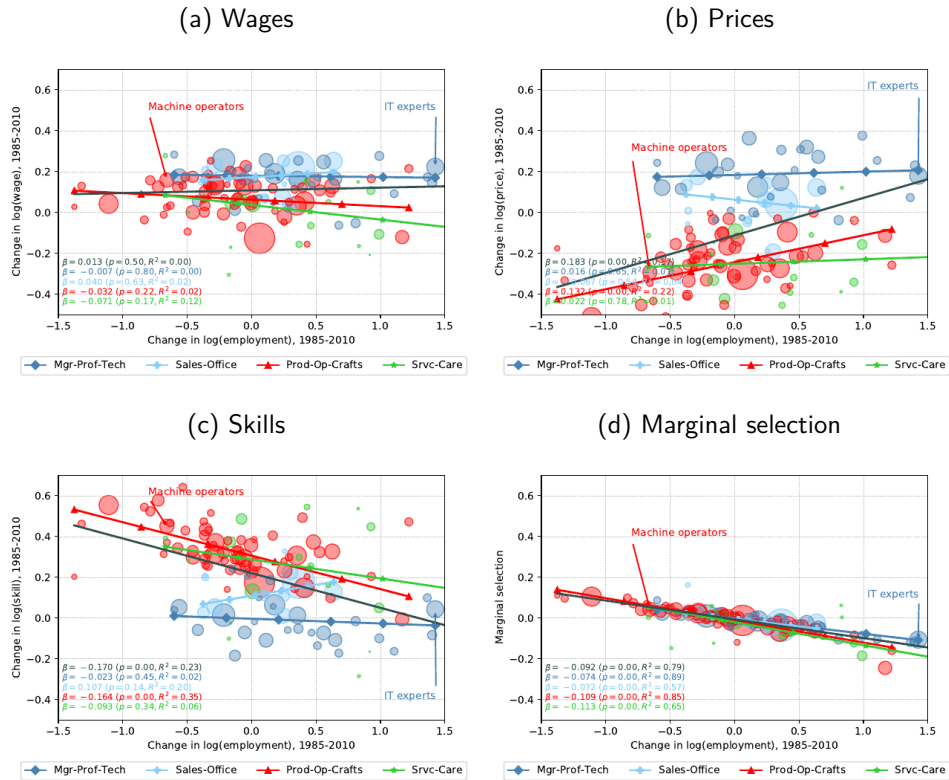
In the main estimation, skill accumulation varies by combination of current and last year's occupation as well as by age in order to account for the differential life-cycle wage growth in these dimensions. In Figure 1.46, we also allow for the fact that skill accumulation may additionally vary by the worker's education level on top of detailed occupation and age. Practically, considering skill accumulation Equation (1.7), we add dummies for high (university or college degree), medium (apprenticeship or Abitur), and low (without postsecondary) education level to the worker characteristics vector $X_{i,t-1}$ and the according coefficients to the $\Gamma_{k',k}$ skill accumulation parameter matrix. Skill accumulation is faster for highly educated workers in almost every occupation. Nonetheless, the elasticity of skill price changes with respect to employment growth is only slightly lower than the baseline elasticity, and the other results are also similar.

The next specification allows for changing relative average occupation-specific amenities by age group over time as described in Section 1.B.3. In particular, we implement Equation (1.28), adding the change in choices $\Delta I_{k,i,t}$ between two periods as a regressor (interacted with $X_{i,t-1}$) to the estimation. We perform this exercise for the four broad occupations only because of the extensive data requirements. Figure 1.47 plots the resulting skill prices and skill accumulation coefficients of the four broad occupation groups, showing that they are hardly affected by this augmented estimation model compared to the main text. In addition, we can also identify the changes of amenities themselves in this specification. Panels 1.47d to 1.47f show the amenities relative to the omitted Prod-Op-Crafts and the base period by age group. We see that for 45–55 year olds these are about zero and pretty stable over time. For young 25–34 olds we do see declining amenities (i.e., rising estimation coefficient on $\Delta I_{k,i,t}$), first for Mgr-Prof-Tech and Sales-Office after the mid-1990s and then for Srvc-Care in the early 2000s. The middle-aged workers are somewhere in between old and young ages with possibly a slight decline of amenities in the other three occupation groups compared to Prod-Op-Crafts toward the end of the sample period.

Finally, in Figure 1.48 we also compare our results to the alternative estimation method using fixed effects due to Cortes (2016). As discussed in Appendix 1.B.4, in order not to control for worker's entire labor market history, we implement it with individual fixed effects for each occupation stint (as do Cavaglia and Etheridge, 2017). With large idiosyncratic skill shocks, there can be some bias in the fixed effects approach, which is illustrated in the Monte Carlo simulations of Appendix 1.C.6. Nonetheless, it

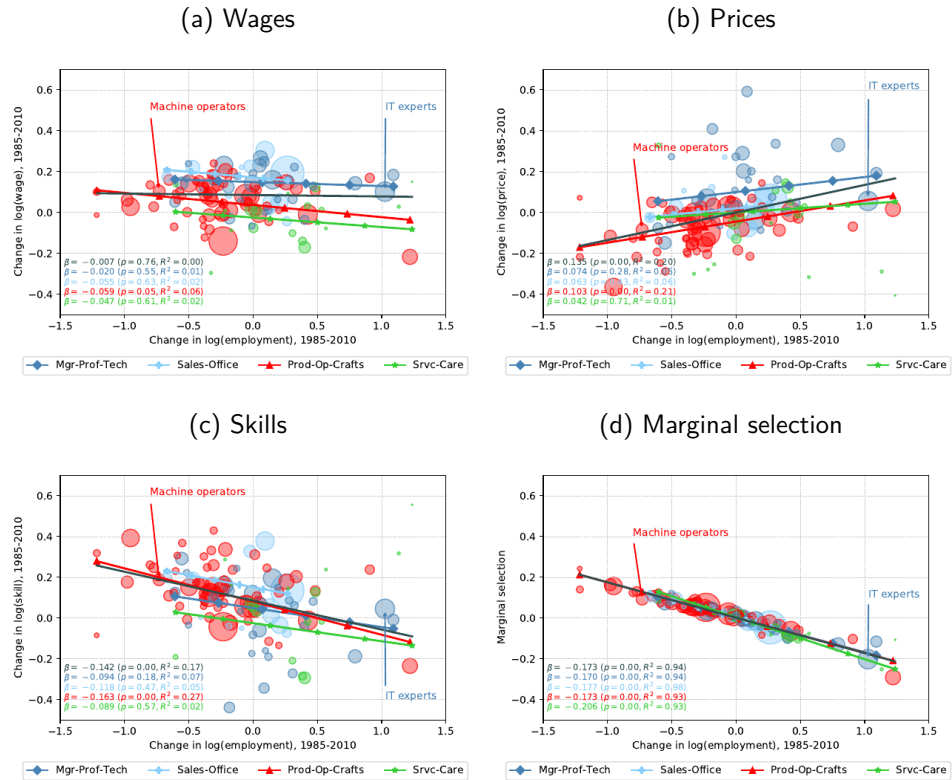
should be supportive of our empirical results that this alternative estimation method yields qualitatively similar findings. For example, the elasticity of skill price changes with respect to employment changes is 0.11 compared to the slope in our baseline method of 0.15.

Figure 1.45: Flat spot identification using workers aged 45–54 years



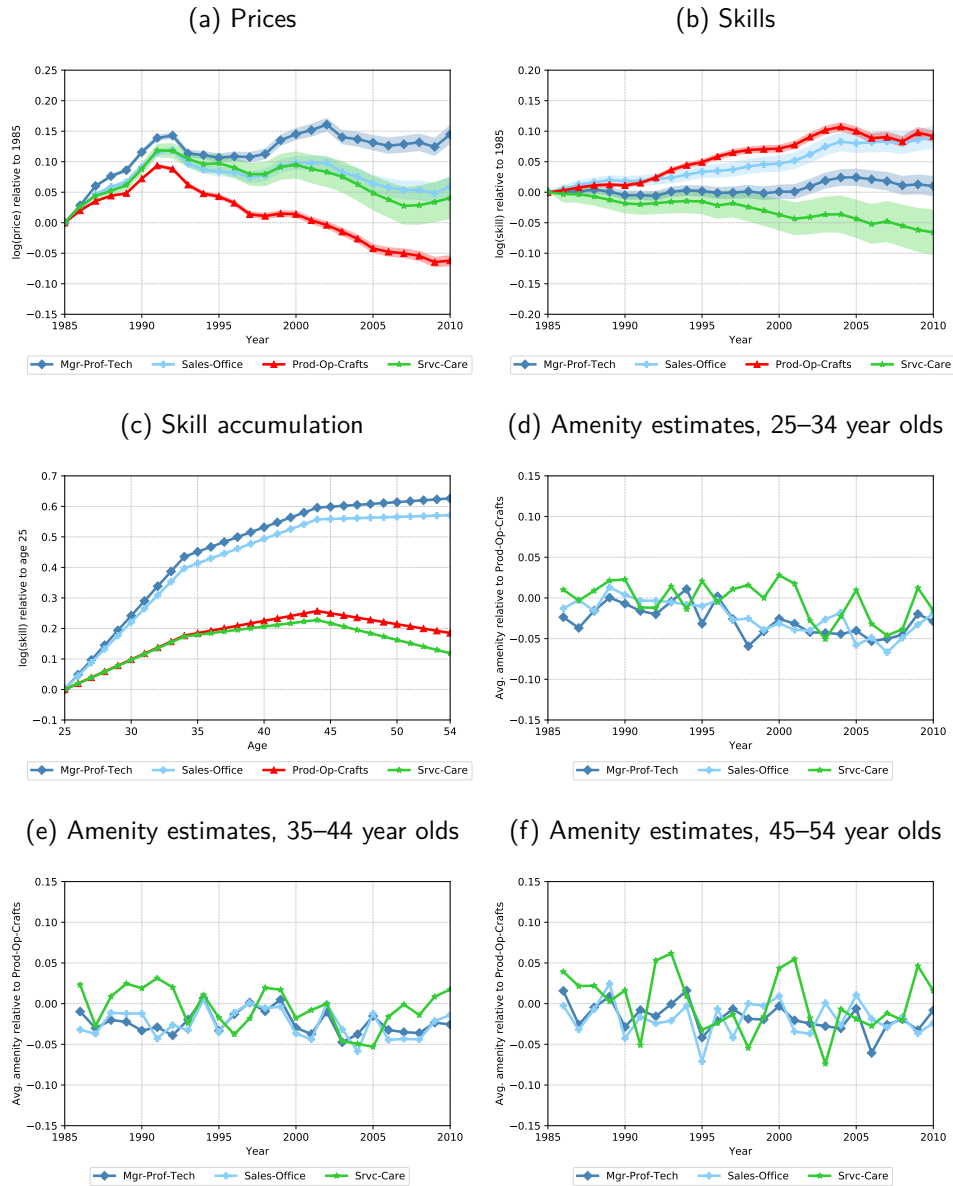
Notes: The sample is restricted to 45–54 year old men. Skill accumulation is set to zero across occupations. The horizontal axes in all panels show the change of the log number of employed workers within an occupation between 1985 and 2010. One bubble represents one of the 120 detailed occupations in the SIAB SUF. The four groups show an aggregation of these detailed occupations as described in Appendix Table 1.4. Bubble size corresponds to the number of employed workers in an occupation averaged across years 1985 until 2010. Regression lines across all occupations (black) and within the four broad groups (colored) are weighted by the number of employed workers.

Figure 1.46: Education-specific skill accumulation



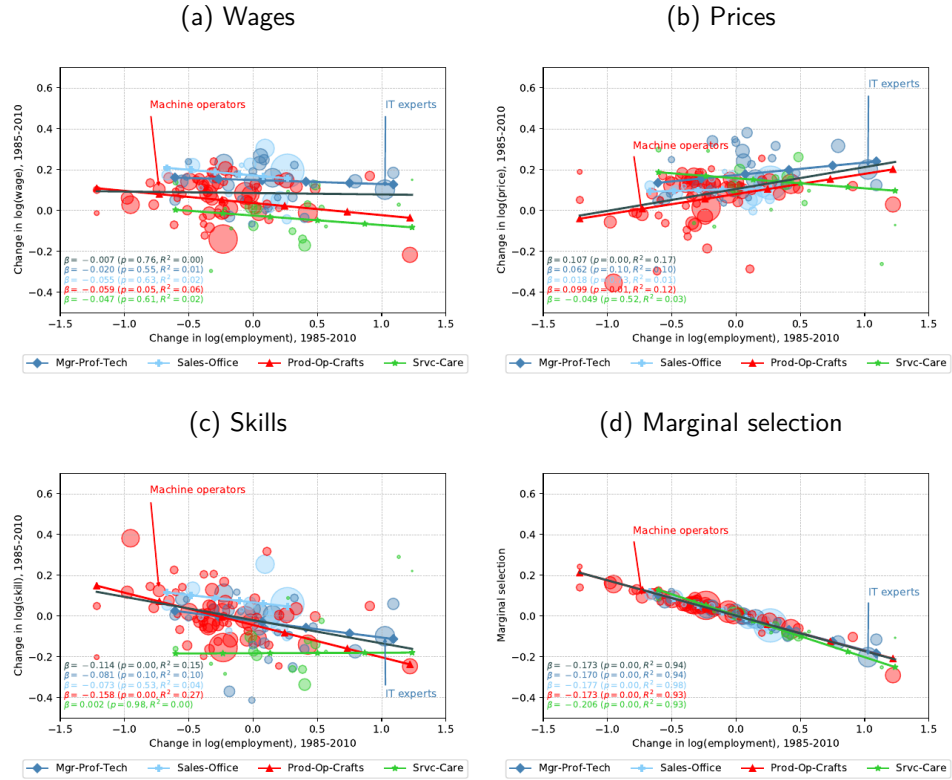
Notes: The speed of skill accumulation described by Equation (1.7) is allowed to vary with worker's education by including dummies for three education levels (low (missing or without any postsecondary education), medium (apprenticeship training or high school diploma), and high (university degree)). The horizontal axes in all panels show the change of the log number of employed workers within an occupation between 1985 and 2010. One bubble represents one of the 120 detailed occupations in the SIAB SUF. The four groups show an aggregation of these detailed occupations as described in Appendix Table 1.4. Bubble size corresponds to the number of employed workers in an occupation averaged across years 1985 until 2010. Regression lines across all occupations (black) and within the four broad groups (colored) are weighted by the number of employed workers.

Figure 1.47: Accounting for non-pecuniary benefits



Notes: Panels 1.47d to 1.47f present the results the three age groups contained in the main sample. Estimates from the extended model which allows to control for changing amenity values of occupations as described by Equation (1.28). Amenities are allowed to vary by age group and identified relative to Prod-Op-Crafts.

Figure 1.48: Stint fixed effects estimation



Notes: Estimates are identified from year specific occupation fixed effects while including a separate worker-occupation fixed effects for each time the worker revisits an occupation (after a possible break or after return from another occupation). Additionally, we include controls for age and occupation dependent skill accumulation following Equation (1.39). The horizontal axes in all panels show the change of the log number of employed workers within an occupation between 1985 and 2010. One bubble represents one of the 120 detailed occupations in the SIAB SUF. The four groups show an aggregation of these detailed occupations as described in Appendix Table 1.4. Bubble size corresponds to the number of employed workers in an occupation averaged across years 1985 until 2010. Regression lines across all occupations (black) and within the four broad groups (colored) are weighted by the number of employed workers.

1.F.3 Task Measures and Changes in Occupations' Employment, Wages, Prices, and Skills

Our data on tasks that workers typically perform on the job comes from the Qualification and Career Surveys (QCS), which are conducted by the Federal Institute for Vocational Education and Training (BIBB) and have previously been used to study task intensities (see for instance Spitz-Oener (2006), Antonczyk et al. (2009), Black and Spitz-Oener (2010), or Gathmann and Schönberg (2010)). The QCS are representative cross-sectional surveys with roughly 20000–35000 respondents in each wave. There are six waves available, which were conducted in 1979, 1985/86, 1991/92, 1998/99, 2006 and 2012. The surveys contain detailed questions about tasks that are required in the workers' occupations, such as how often they repair objects or how often they have to persuade co-workers. We classify each question as representing either analytical, interactive, routine, or manual tasks and assign a value of 0, $\frac{1}{3}$ or 1, depending on whether the answer is 'never', 'sometimes', or 'frequently' (or 0/1 for yes/no questions). Since the questions are not always comparable across waves, we pool all waves to compute task intensities by averaging over all responses. Note that the intensities are constructed in a way so that the four dimensions do not sum to 1, which follows the approach in Spitz-Oener (2006). There are two types of variation in responses that lead to variation in absolute task intensities across occupations. First, at the "extensive margin" fewer or more workers can reply that they engage in a task that is asked for in a specific question at all. As an example, consider the simple case with two questions about analytical tasks and individuals in one occupation doing both tasks and individuals in the other occupation only one. Second, at the "intensive margin", individuals in the occupation could more or less often reply that they engage 'sometimes' in a task, as opposed to 'frequently'.

1.G Further Details on Wage Inequality

This section provides more details of the analyses of the wage distribution. We begin with the inequality between occupations from Section 1.5.1. Then we turn to further analyses and robustness checks for the effect of the full estimated model on various wage percentiles and overall inequality.

1.G.1 Derivations and Further Results on the Attenuating Effect of Selection

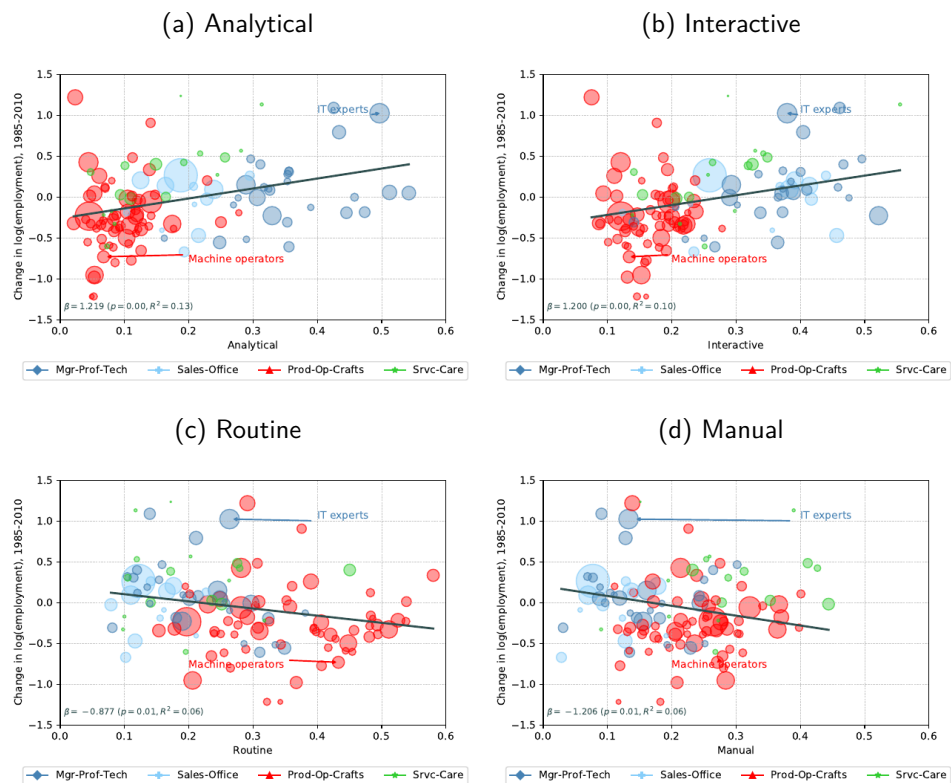
We start by deriving the Decomposition (1.14) of the main text. Note that $\sigma(\bar{w}_{k,t}, \bar{w}_{k,1985}) = cov(\bar{w}_{k,t}, \bar{w}_{k,1985})$, $\sigma^2(\bar{w}_{k,t}) = \sigma(\bar{w}_{k,t}, \bar{w}_{k,t})$, and by the linear additivity of the covariance operator:

$$\sigma^2(\Delta\bar{w}_{k,t}) = \sigma^2(\bar{w}_{k,t} - \bar{w}_{k,1985}) = \sigma^2(\bar{w}_{k,t}) + \sigma^2(\bar{w}_{k,1985}) - 2 \cdot \sigma(\bar{w}_{k,t}, \bar{w}_{k,1985})$$

Rearranging this give the terms under the braces in Equation (1.14):

$$\begin{aligned} \Delta\sigma^2(\bar{w}_{k,t}) &= \sigma^2(\Delta\bar{w}_{k,t}) - 2 \cdot \sigma^2(\bar{w}_{k,1985}) + 2 \cdot \sigma(\bar{w}_{k,t}, \bar{w}_{k,1985}) \\ &= \sigma^2(\Delta\bar{w}_{k,t}) - 2 \cdot \sigma^2(\bar{w}_{k,1985}) + 2 \cdot \sigma(\bar{w}_{k,1985} + \Delta\bar{w}_{k,t}, \bar{w}_{k,1985}) \quad (1.52) \\ &= \sigma^2(\Delta\bar{w}_{k,t}) + 2 \cdot \sigma(\Delta\bar{w}_{k,t}, \bar{w}_{k,1985}) \end{aligned}$$

Figure 1.49: Correlation of employment changes with task measures



Notes: The vertical axes in all panels show the change of the log number of employed workers within an occupation between 1985 and 2010. Task measures were computed using the Qualifications and Career Surveys. In the QCS surveys, workers are asked what tasks they perform in their job, e.g. “how often do you repair stuff”. They provide answers on a scale: “never, sometimes, often”. We assign numerical values $\{0, \frac{1}{3}, 1\}$ to these categories, respectively. We group all the questions into the four categories mentioned in the headers and average over occupations implying that the four task categories do not need to sum up to one as some occupations might be more intense in overall tasks than others. The six different QCS waves were pooled together as the questions are hardly comparable between waves. One bubble represents one of the 120 detailed occupations in the SIAB SUF. The four groups show an aggregation of these detailed occupations as described in Appendix Table 1.4. Bubble size corresponds to the number of employed workers in an occupation averaged across years 1985 until 2010. Regression lines across all occupations (black) and within the four broad groups (colored) are weighted by the number of employed workers.

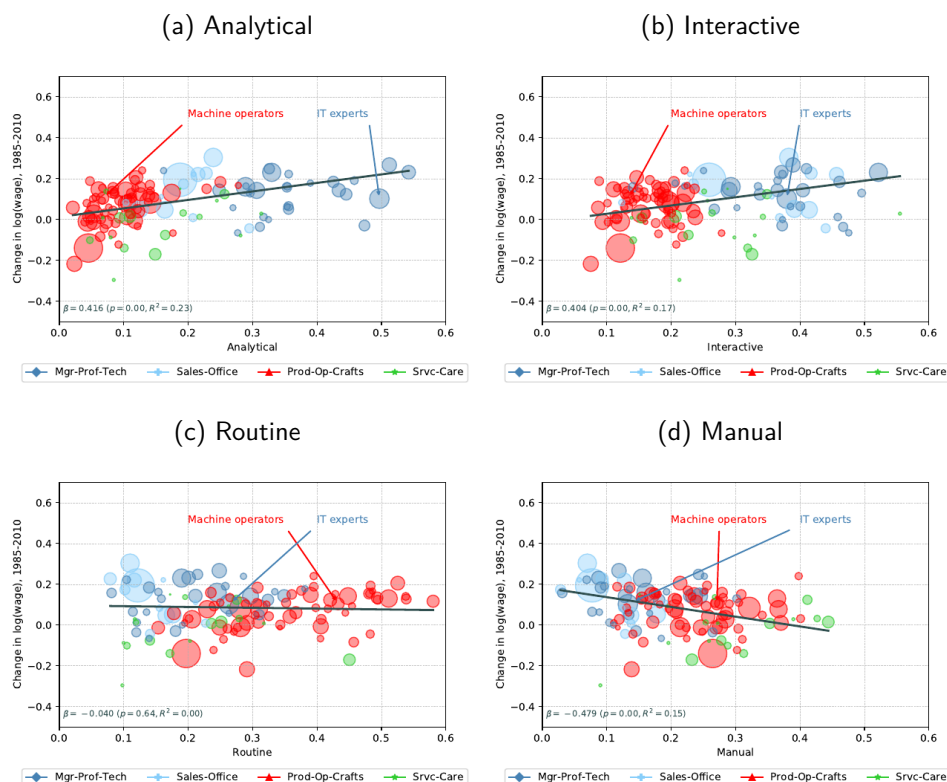
Finally, inserting the sum of skill prices and average skills for the average wages (i.e., $\bar{w}_{k,t} = \pi_{k,t} + \bar{s}_{k,t}$):

$$\begin{aligned} \Delta\sigma^2(\bar{w}_{k,t}) &= \sigma^2(\Delta\pi_{k,t} + \Delta\bar{s}_{k,t}) + 2 \cdot \sigma(\Delta\pi_{k,t} + \Delta\bar{s}_{k,t}, \bar{w}_{k,1985}) \\ &= \sigma^2(\Delta\pi_{k,t}) + \sigma^2(\Delta\bar{s}_{k,t}) + 2 \cdot \sigma(\Delta\pi_{k,t}, \Delta\bar{s}_{k,t}) \\ &\quad + 2 \cdot \sigma(\Delta\pi_{k,t}, \bar{w}_{k,1985}) + 2 \cdot \sigma(\Delta\bar{s}_{k,t}, \bar{w}_{k,1985}) \end{aligned} \quad (1.53)$$

where $\bar{s}_{k,t}$ is the average skill in occupation k at time t .

The second panel of Table 1.24 restates the results of this decomposition, in the actual data and for the counterfactuals with price changes only, reweighting of the demographic structure, and the combination of the two. This is all the same as in the main

Figure 1.50: Correlation of wage changes with task measures

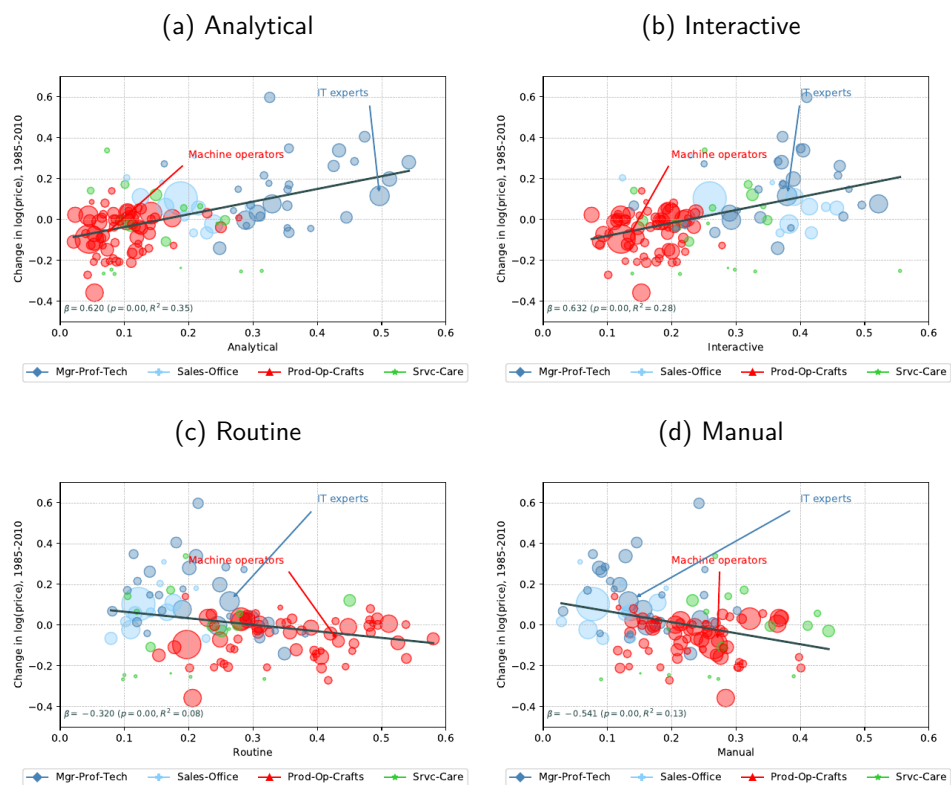


Notes: The vertical axes in all panels show the change of the average log wage within an occupation between 1985 and 2010. Task measures were computed using the Qualifications and Career Surveys. In the QCS surveys, workers are asked what tasks they perform in their job, e.g. “how often do you repair stuff”. They provide answers on a scale: “never, sometimes, often”. We assign numerical values $\{0, \frac{1}{3}, 1\}$ to these categories, respectively. We group all the questions into the four categories mentioned in the headers and average over occupations implying that the four task categories do not need to sum up to one as some occupations might be more intense in overall tasks than others. The six different QCS waves were pooled together as the questions are hardly comparable between waves. One bubble represents one of the 120 detailed occupations in the SIAB SUF. The four groups show an aggregation of these detailed occupations as described in Appendix Table 1.4. Bubble size corresponds to the number of employed workers in an occupation averaged across years 1985 until 2010. Regression lines across all occupations (black) and within the four broad groups (colored) are weighted by the number of employed workers.

text, other than that—at the very right of the table—we add another counterfactual where also the occupation distribution is used for the reweighting. That counterfactual is not much different from the one before where only age, foreigner status, and education are reweighted; as we have also mentioned in the main text.

Within-occupation wage inequality may also be affected by selection into occupations due to changing skill prices. If, for instance, rising prices attract workers of lower skill than the incumbents, inequality will increase within growing sectors. If occupations with high inequality grow, then within inequality will rise overall. Conversely, within inequality might decrease in shrinking occupations with declining skill prices because their low skilled workers may leave. The bottom panel of Table 1.24 also shows a

Figure 1.51: Correlation of skill price changes with task measures

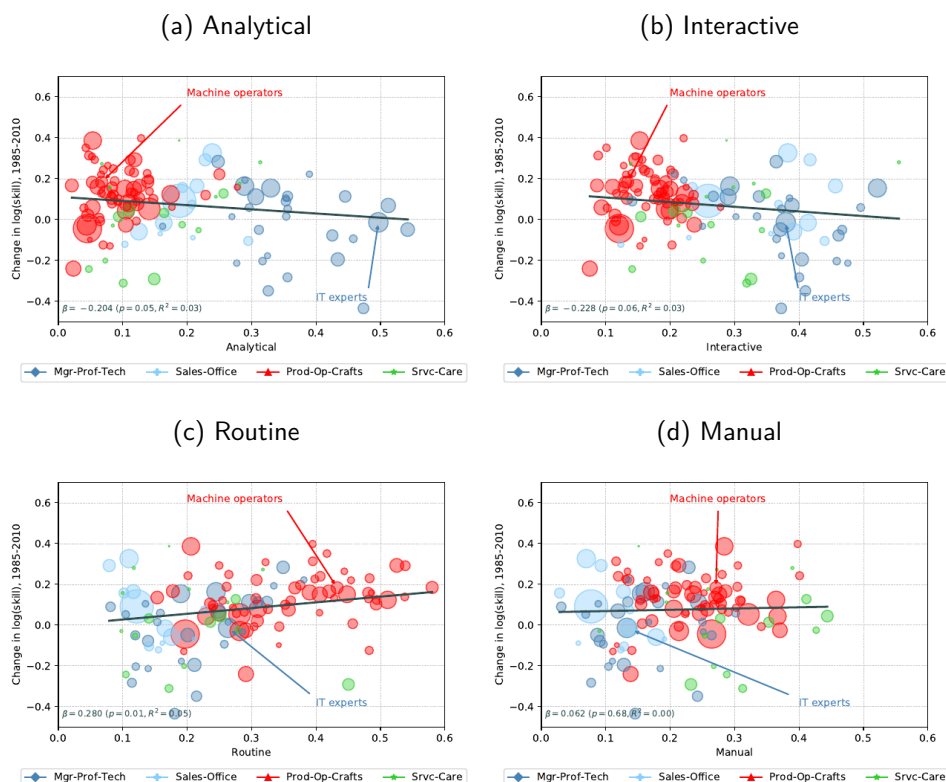


Notes: The vertical axes in all panels show the change of skill prices between 1985 and 2010. OLS estimates as described by Equation (1.9). Task measures were computed using the Qualifications and Career Surveys. In the QCS surveys, workers are asked what tasks they perform in their job, e.g. “how often do you repair stuff”. They provide answers on a scale: “never, sometimes, often”. We assign numerical values $\{0, \frac{1}{3}, 1\}$ to these categories, respectively. We group all the questions into the four categories mentioned in the headers and average over occupations implying that the four task categories do not need to sum up to one as some occupations might be more intense in overall tasks than others. The six different QCS waves were pooled together as the questions are hardly comparable between waves. One bubble represents one of the 120 detailed occupations in the SIAB SUF. The four groups show an aggregation of these detailed occupations as described in Appendix Table 1.4. Bubble size corresponds to the number of employed workers in an occupation averaged across years 1985 until 2010. Regression lines across all occupations (black) and within the four broad groups (colored) are weighted by the number of employed workers.

decomposition of wage inequality within occupations.

Denote $\tilde{w}_{k,i,t}$ as the difference between an individual's wage and the average wage within his occupation. Given that skill prices are the same for a fixed occupation, this residual wage difference is the same as the residual skill difference: $\tilde{w}_{k,i,t} = \tilde{s}_{k,i,t} =$

Figure 1.52: Correlation of skill changes with task measures



Notes: The vertical axes in all panels show the change of skills between 1985 and 2010 estimated as the residual between price and wage changes as shown in Equation (1.12). OLS estimates as described by Equation (1.9). Task measures were computed using the Qualifications and Career Surveys. In the QCS surveys, workers are asked what tasks they perform in their job, e.g. “how often do you repair stuff”. They provide answers on a scale: “never, sometimes, often”. We assign numerical values $\{0, \frac{1}{3}, 1\}$ to these categories, respectively. We group all the questions into the four categories mentioned in the headers and average over occupations implying that the four task categories do not need to sum up to one as some occupations might be more intense in overall tasks than others. The six different QCS waves were pooled together as the questions are hardly comparable between waves. One bubble represents one of the 120 detailed occupations in the SIAB SUF. The four groups show an aggregation of these detailed occupations as described in Appendix Table 1.4. Bubble size corresponds to the number of employed workers in an occupation averaged across years 1985 until 2010. Regression lines across all occupations (black) and within the four broad groups (colored) are weighted by the number of employed workers.

$s_{k,i,t} - \bar{s}_{k,t}$. The average within-occupation variance of log wages becomes:

$$\sigma^2(\tilde{w}_{k,i,t}) = \sigma^2(\tilde{s}_{k,i,t}) = \frac{1}{N_t} \sum_{i=1}^{N_t} \tilde{s}_{k,i,t}^2 = \sum_{k=1}^K \frac{p_{k,t}}{N_t} \cdot \overbrace{\frac{1}{N_{k,t}} \sum_{i \in k} \tilde{s}_{k,i,t}^2}^{\sigma_{k,t}^2} = \sum_{k=1}^K p_{k,t} \sigma_{k,t}^2, \quad (1.54)$$

where $p_{k,t}$ is occupation k 's share of total employment at time t and $\sigma_{k,t}^2$ is the variance

of wages or skills within the occupation. The change of the average within variance is:

$$\begin{aligned}\Delta\sigma^2(\tilde{w}_{k,i,t}) &= \Delta\sigma^2(\tilde{s}_{k,i,t}) = \sum_{k=1}^K (p_{k,t}\sigma_{k,t}^2 - p_{k,1985}\sigma_{k,1985}^2) \\ &= \sum_{k=1}^K \Delta\sigma_{k,t}^2 p_{k,1985} + \sum_{k=1}^K \Delta p_{k,t} \sigma_{k,1985}^2 + \sum_{k=1}^K \Delta\sigma_{k,t}^2 \Delta p_{k,t}\end{aligned}\tag{1.55}$$

Therefore, the rise of within-occupation inequality can be decomposed into terms linked to the changing employment structure and 'pure' increases of the variance of log wages in occupations at fixed sizes. In addition, the last summand of Equation (1.55) is actually the covariance of changing within inequality with changing employment share. That is, how much the variance of skills in occupations rises for growing occupations, which is related to the declining (rising) skills in growing (shrinking) occupations discussed in Section 1.4. This relationship generates 0.56 (i.e., $1.34 - 0.78$) log points of the increase in within inequality in the second column, bottom panel of Table 1.24.

The other component related to the changing employment structure is the growing size of sectors with high initial within inequality. These are often relatively large occupations inside the rising Mgr-Prof-Tech, Sales-Office, and Srv-Care groups, which is partly due to the German KLDB occupation classification as it is more detailed in production and crafts related occupations than in managerial, office, or service type occupations (see Table 1.4). The effect of this is the second summand in Equation (1.55) and it makes up another 0.59 log points of the increase in within-occupation inequality in Table 1.24. Clearly the largest part of the rising within variance is the first summand in Equation (1.55). However, also here the employment structure played a role because larger occupations, which as we said are often in Mgr-Prof-Tech, Sales-Office, and Srv-Care occupations, had higher increases of within inequality.

The remaining columns in the bottom panel of Table 1.24 again show and decompose the effect of the reweighting counterfactuals (notice that the skill prices vary only across occupations and thus have no effect on inequality within). We see that this has an overall modest effect but that additionally reweighting the occupation composition at the very right of the table does raise that effect, since it almost perfectly matches the growth of occupations with large inequality within (i.e., 0.59 in the actual and 0.63 in that counterfactual).

Table 1.24: Between-within occupation variance of observed log wages, experiments

		Level 1985		Difference 2010 - 1985						
		Observed	Observed	Prices only	Rew. age, foreign	+ prices	Rew. age, for- eign, educa- tion	+ prices	Rew. age, for- eign, educa- tion, occu- pation	+ prices
Overall	$\sigma^2(w_{i,t})$	14.28	12.41	5.13	0.45	5.70	2.41	9.56	2.70	8.44
Between	$\sigma^2(\bar{w}_{k,t})$	5.03	5.25	5.13	0.34	5.59	1.74	8.88	1.60	7.34
	$2 \cdot \sigma(\Delta\pi_{k,t}, \bar{w}_{k,1985})$		3.23	3.23	0.00	3.23	0.00	3.23	0.00	3.23
	$\sigma^2(\Delta\pi_{k,t})$		1.76	1.89	0.00	1.92	0.00	2.36	0.00	1.79
	$2 \cdot \sigma(\Delta\bar{s}_{k,t}, \bar{w}_{k,1985})$		-0.53	0.00	0.26	0.26	0.53	0.53	0.59	0.59
	$\sigma^2(\Delta\bar{s}_{k,t})$		3.02	0.00	0.07	0.07	1.21	1.21	1.00	1.00
	$2 \cdot \sigma(\Delta\pi_{k,t}, \Delta\bar{s}_{k,t})$		-2.24	0.00	0.00	0.10	0.00	1.55	0.00	0.73
Within	$\sigma^2(w_{i,t} - \bar{w}_{k,t})$	9.25	7.16	0.00	0.11	0.11	0.68	0.68	1.10	1.10
	$\sum_k \Delta\sigma_{k,t}^2 \Delta p_{k,t}^{>0}$		1.34	0.00	0.00	0.00	-0.01	-0.01	0.05	0.05
	$\sum_k \Delta\sigma_{k,t}^2 \Delta p_{k,t}^{\leq 0}$		-0.78	0.00	-0.00	-0.00	-0.03	-0.03	-0.07	-0.07
	$\sum_k \Delta p_{k,t} \sigma_{k,1985}^2$		0.59	0.00	-0.01	-0.01	0.28	0.28	0.63	0.63
	$\sum_k \Delta\sigma_{k,t}^2 p_{k,1985}$		6.02	-0.00	0.13	0.13	0.44	0.44	0.49	0.49

Notes: The levels in 1985 are 14.3 (overall), 5.0 (between), and 9.3 (within). Based on specification with 120 occupations. $\bar{w}_{k,t}$ refers to the average wage in occupation k in year t . Values $\times 100$. The experiments are: Prices only: Take average wages in 1985 and add price changes to simulate 2010's wages. Rew. age, foreign: Take wages in 1985 and reweigh them according to 2010's age and foreign worker distribution with weights computed following DiNardo et al. (1996) to simulate 2010's wages. + prices: add price changes on top.

1.G.2 Additional Results for the Scenarios from the Full Model

Here, we first provide the explicit formulas for the scenarios in Figure 1.9. We then describe some features of the table and figures in this section. These include the levels of the percentiles and the variance in the data and under the model prediction (Table 1.25). They also include two different sample/data preparation specifications (Figures 1.53 and 1.54), and a permutation of the order in which we add the components of the model (Figure 1.55).

The first scenario in Figure 1.9c reports the trend of inequality that would prevail if only the wage distribution at age 25 (or an older age for later entrants) had shifted, with changes in skill prices as well as skill accumulation or occupation-switching over the life-cycle turned off. That is,

$$\hat{w}_{i,t}^{\text{Initial occupation and wage throughout}} = w_{i,t_{i,0}}, \quad (1.56)$$

where $t_{i,0} \leq t$ denotes the year when worker i joins the labor market.

The next scenario adds skill accumulation to these initial wages. In particular, Figure 1.9d shows the inequality due to changing initial occupation distribution (measured by $I_{k,i,t_{i,0}}$), age structure of employment ($X_{i,\tau-1}$), and associated changes of skill accumulation over time ($\hat{\Gamma}_{k,k}$):

$$\hat{w}_{i,t}^{\text{Initial occupation + skill accumulation}} = w_{i,t_{i,0}} + \sum_{\tau=t_{i,0}+1}^t \sum_{k=1}^K I_{k,i,t_{i,0}} I_{i,\tau-1} \cdot X'_{i,\tau-1} \hat{\Gamma}_{k,k}, \quad (1.57)$$

where the worker joined the labor market at time $t_{i,0} \leq t$, never switches (i.e., $I_{k,i,t_{i,0}}$ does not change with τ), and $I_{i,\tau-1}$ indicates whether the worker was employed in $\tau - 1$. That is, we assume that skills stagnate during non-employment spells. Also, workers who are currently unemployed or out of the labor force do not enter any of the scenarios.

The scenario in Figure 1.9e includes occupation switching, but leaves out the skill changes associated with switching. The wage in this scenario becomes:

$$\hat{w}_{i,t}^{\text{Observed occ. + skill acc.; } \Delta s_{k,t} = 0} = w_{i,t_{i,0}} + \sum_{\tau=t_{i,0}+1}^t \sum_{k=1}^K I_{k,i,\tau-1} \cdot X'_{i,\tau-1} \hat{\Gamma}_{k,k} \quad (1.58)$$

Here notably $I_{k',i,\tau-1}$ can be zero for all K occupations if the worker was unemployed or out of the labor force in the respective period. We do assign workers who return from non-employment one $\hat{\Gamma}_{k,k}$ for their previous occupation, however. In the scenarios with the gains from switching immediately below we assign $\hat{\Gamma}_{k',k}$ for their previous k' and current k occupation combination. Either of this does not make a material difference for the results and, for ease of notation, it is not explicitly indicated in the formulas.

The next scenario adds the skill changes associated with switching in Figure 1.9f. The wage in this scenario becomes:

$$\hat{w}_{i,t}^{\text{Observed occ. + skill accumulation}} = w_{i,t_{i,0}} + \sum_{\tau=t_{i,0}+1}^t \sum_{k=1}^K \sum_{k'=1}^K I_{k,i,\tau} \left(I_{k',i,\tau-1} \cdot X'_{i,\tau-1} \hat{\Gamma}_{k',k} \right) \quad (1.59)$$

Finally, we add our estimated skill prices in Figure 1.9b. The wage in the full empirical model becomes:

$$\begin{aligned} \hat{w}_{i,t}^{\text{Model}} &= w_{i,t_{i,0}} + \sum_{k=1}^K I_{k,i,t} \sum_{\tau=t_{i,0}+1}^t \Delta \hat{\pi}_{k,\tau} \\ &+ \sum_{\tau=t_{i,0}+1}^t \sum_{k=1}^K \sum_{k'=1}^K I_{k,i,\tau} \left(I_{k',i,\tau-1} \cdot X'_{i,\tau-1} \hat{\Gamma}_{k',k} \right) \end{aligned} \quad (1.60)$$

As detailed in the main text, we consider two different specifications in this section. First, to partly abstract from the supply shock of increased migration after 1990, we exclude anybody ever excluded as a foreigner. Second, we fill non-employment spells to see how periods in unemployment or outside of the labor force impact the scenarios. For the latter, all notes related to these spells in the formulas (1.56)–(1.60) are irrelevant here.

Since Figure 1.9 showed the evolution of the three percentiles under consideration normalized to zero in 1985, we display their 1985 levels in Table 1.25 along with the variance. Across all scenarios, the model fits the levels well. None of the values deviates by more than 1.5/100 from the actual value. This is quite remarkable given that the estimation of the model targets average occupational wages for demographic groups. The predictions for the differences between the endpoints of our study period are also broadly in line with the data.

Figure 1.53 shows the same scenarios as Figure 1.9 in the main text when excluding anybody who was ever coded as a foreigner from the sample. The most important fact to note is that in Panel c, the decline in the 15th percentile is not nearly as pronounced as in Figure 1.9c. In the experiment where unemployment is a choice (Figure 1.54), the same broad conclusions hold as in the main text. Some differences are noteworthy, however. Already in the data, the 15th percentile decreases much more than in Figure 1.9a, the pattern is similar but much less pronounced for the higher percentiles. In this specification, price changes hurt both the median and the 15th percentile; the latter actually slightly loses from switching (Panels e → f). This highlights that we overestimate the gains from switching at the lower end because occupation changes involving wage losses often go via an unemployment spell.

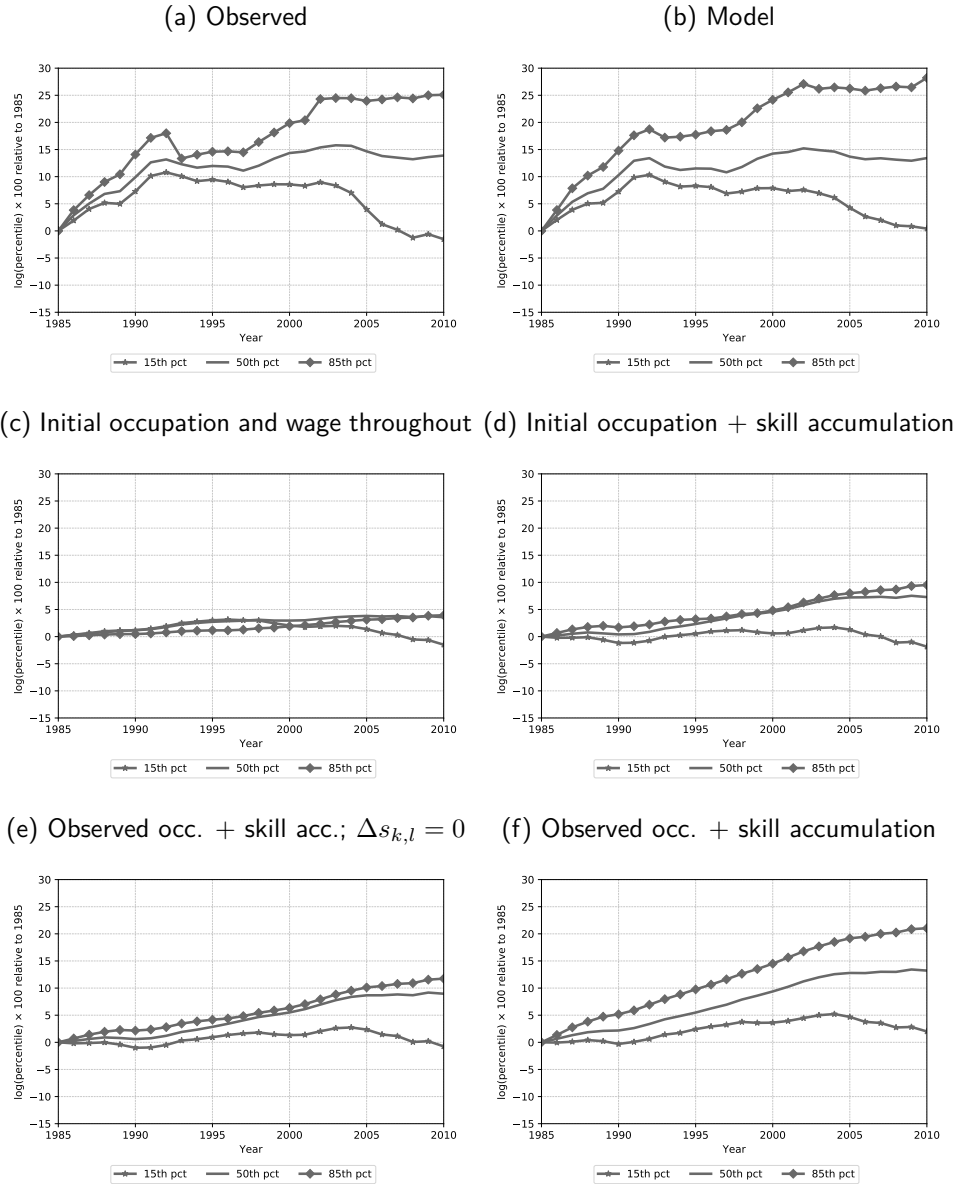
Finally, note that the sequencing of the experiments in Figure 1.9 is arbitrary. In fact, the sequencing only matters with respect to when occupation switching is added because, conditional on occupation choice, all other components enter separately and do not interact. In Figure 1.55, we add prices immediately to the initial wages and add occupation switching as late as possible. Interestingly, all three percentiles would have evolved in almost the same way until the mid-nineties and they opened up only afterwards. Adding skill accumulation then yields the more familiar pattern, which is opened up further by switches.

Table 1.25: Levels of wage percentiles and the variance for model and sample specifications

		Main sample		Anybody ever coded as foreign excluded		Filled non-employment spells	
		Level 1985	Diff. 2010-85	Level 1985	Diff. 2010-85	Level 1985	Diff. 2010-85
High $\log(p_{85})$	Data	1080.71	24.03	1081.42	25.09	1079.12	23.27
	Model prediction	1082.38	26.86	1082.73	28.21	1080.00	23.27
Median $\log(p_{50})$	Data	1042.88	11.55	1043.14	13.89	1041.27	9.75
	Model prediction	1044.48	11.63	1044.78	13.42	1041.65	5.96
Low $\log(p_{15})$	Data	1017.07	-5.13	1017.78	-1.52	1013.23	-10.53
	Model prediction	1015.95	-2.23	1016.36	0.43	1011.46	-11.51
Variance $\sigma^2(w_{i,t})$	Data	14.28	12.41	14.20	11.21	16.48	15.59
	Model prediction	14.66	12.11	14.63	10.91	16.41	15.22

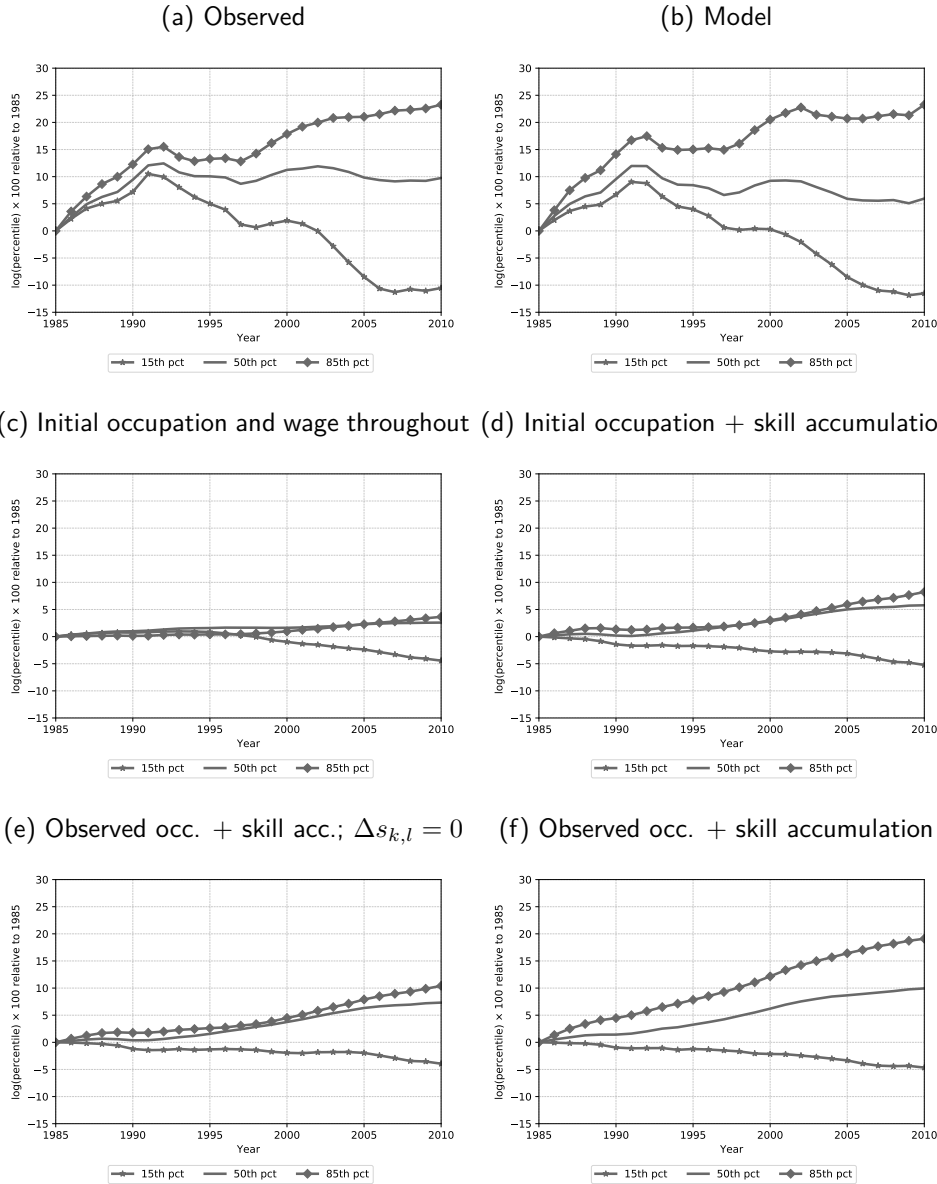
Notes: The table shows observed levels of the 85th, 50th and 15th log wage distribution percentiles as well as the variance of log wages in 1985 and changes between 2010 and 1985. All values $\times 100$. In addition, the table shows model predictions of levels and changes according to Equation (1.60). The first two columns show the results for the baseline sample. Columns three and four present the results when we also drop workers which are reported to be foreigners at some point in time. This includes, for instance, workers acquiring the German nationality at some later point in the life cycle. For the last two columns, unemployment and out of labor force spells are imputed by comparing the (real) wage after a non-employment spell with the wage before the non-employment spell. We then fill up the wage while in non-employment as the lower of those two wages adjusted for inflation and set the occupation within this time to the occupation that corresponds to that lower wage. See Appendix 1.A.1.4 for the details.

Figure 1.53: Wage inequality scenarios, anybody ever coded as foreign excluded

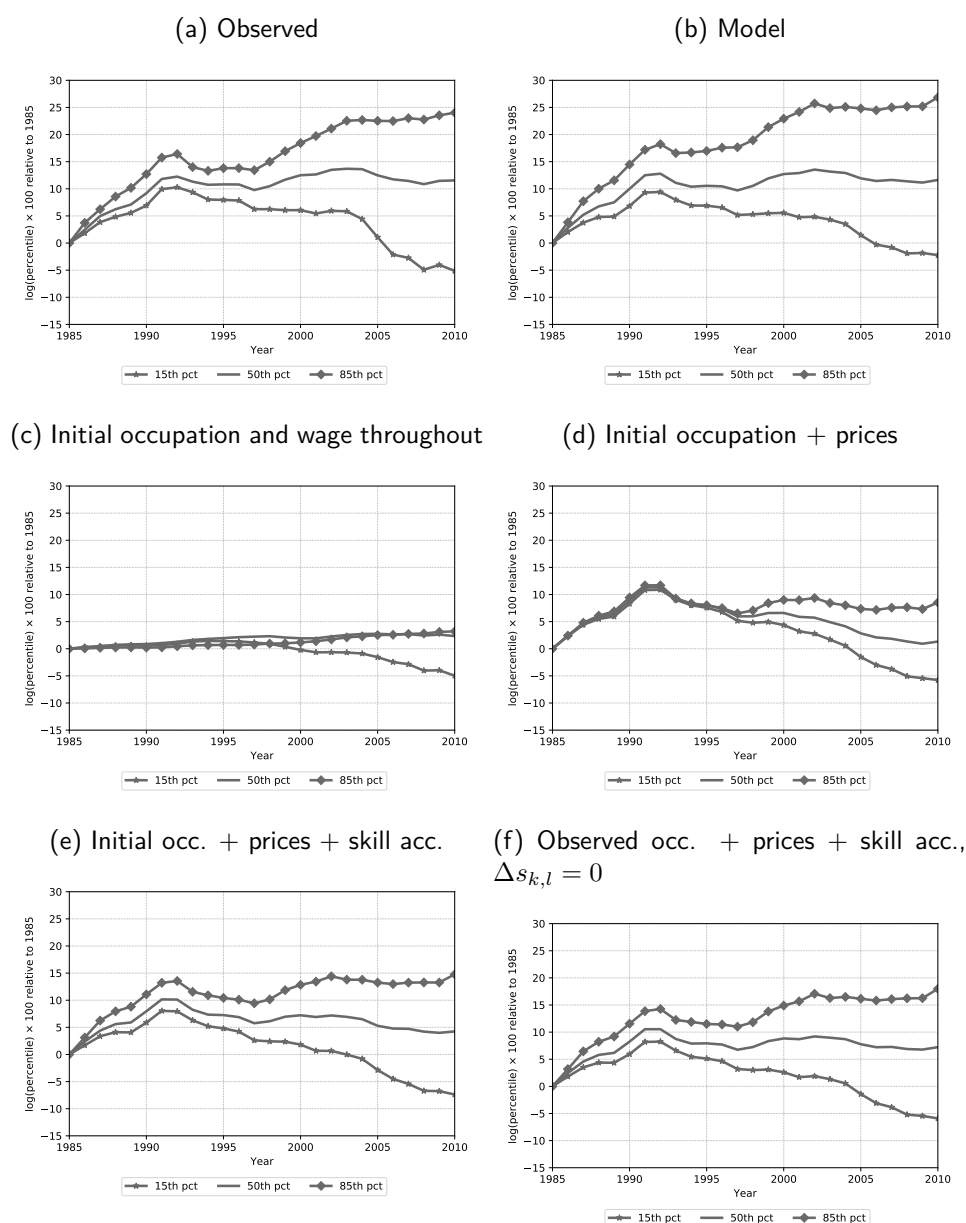


Notes: Panel a: observed wages. Panel b: simulated life-cycle trajectories based on our full model: starting from the initial wage and occupational choice, add all skill accumulation and price change estimates using occupational choices observed in the data. Panel c: workers keep their initial wage throughout the life cycle. Panel d: workers stay in their initial job throughout the life-cycle; in each period, we add the skills they would have accumulated in that job (i.e., $\Delta s_{k_0, k_0, t}$). Panel e: use observed switches, setting direct gains from switching to zero, i.e., $\Delta s_{k,l} = 0 \forall k \neq l$. Price changes are zero as well, so the difference to Panel d comes purely from differential skill accumulation in occupations. Panel f: as in Panel e, but adding the direct gains from switching. The only difference to the full model in Panel b are the price changes, which continue to be zero. In all scenarios, we treat unemployment or out-of-the-labor force spells as follows: When such a spell is observed in the data, simulated workers do not enter the inequality statistics. Furthermore, we assume no depreciation and upon re-entry into paid work add—where relevant—the $\Delta s_{k,l,t}$ with l being the occupation before the spell. The sample is the same as the baseline sample except that we also drop workers which are reported to be foreigners at some point in time. This includes, for instance, workers acquiring the German nationality at some later point in the life cycle.

Figure 1.54: Wage inequality scenarios, filled non-employment spells



Notes: *Panel a*: observed wages. *Panel b*: simulated life-cycle trajectories based on our full model: Starting from the initial wage and occupational choice, add all skill accumulation and price change estimates using occupational choices observed in the data. *Panel c*: workers keep their initial wage throughout the life cycle. *Panel d*: workers stay in their initial job throughout the life-cycle; in each period, we add the skills they would have accumulated in that job (i.e., $\Delta s_{k_0, k_0, t}$). *Panel e*: use observed switches, setting direct gains from switching to zero, i.e., $\Delta s_{k,l} = 0 \forall k \neq l$. Price changes are zero as well, so the difference to Panel d comes purely from differential skill accumulation in occupations. *Panel f*: as in Panel e, but adding the direct gains from switching. The only difference to the full model in Panel b are the price changes, which continue to be zero. Unemployment and out of labor force spells are imputed by comparing the (real) wage after a non-employment spell with the wage before the non-employment spell. We then fill up the wage while in non-employment as the lower of those two wages adjusted for inflation and set the occupation within this time to the occupation that corresponds to that lower wage. See Appendix 1.A.1.4 for the details.

Figure 1.55: Wage inequality scenarios, order prices \rightarrow accumulation \rightarrow switching

Notes: *Panel a:* observed wages. *Panel b:* simulated life-cycle trajectories based on our full model: starting from the initial wage and occupational choice, add all skill accumulation and price change estimates using occupational choices observed in the data. *Panel c:* workers keep their initial wage throughout the life cycle. *Panel d:* workers stay in their initial job throughout the life-cycle; in each period, we add the price changes estimates in that job (i.e., $\Delta\pi_{k_0,t}$). *Panel e:* in addition to Panel d, add differential skill accumulation in occupations. *Panel f:* as in Panel e, but take observed occupational choices as opposed to initial choices. In all scenarios, we treat unemployment or out-of-the-labor force spells as follows: when such a spell is observed in the data, simulated workers do not enter the inequality statistics. Furthermore, we assume no depreciation and upon re-entry into paid work add—where relevant—the $\Delta s_{k,l,t}$ with l being the occupation before the spell.

1.G.3 Effect of Skill Accumulation on Wage Percentiles

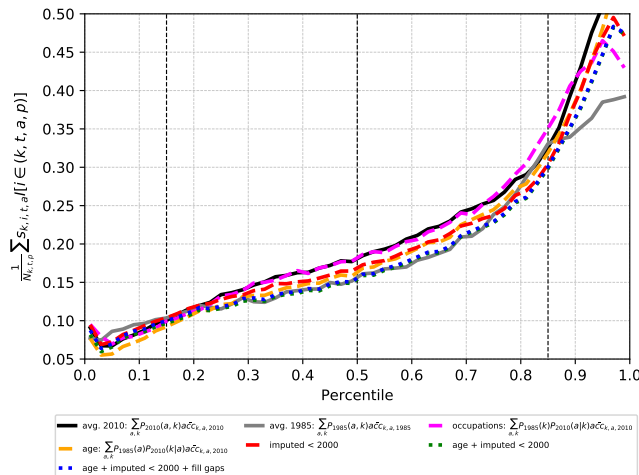
This section analyzes the reasons for skill accumulation’s substantial effect on the change of lower-half inequality, also in comparison to its modest effect on the upper-half. To do this, we rewrite the overall skill accumulation as a function of average skill accumulation within detailed worker cells times the frequency of these cells. In particular, for every percentile of the wage distribution in a given year, we compute the average skill accumulation in each worker cell defined by age and initial occupation. Then we average over these cells by their shares in the respective percentiles. That is, we compute the average skill accumulation in each percentile had workers stayed in the occupation of when they first joined the labor market as:

$$avg_t = \sum_{k=1}^{120} \sum_{a=25}^{54} P_t(a, k) \cdot \overline{acc}_{k,a,t} \quad \text{with} \quad (1.61)$$

$$\overline{acc}_{k,a,t} = \frac{1}{N_{k,a,t}} \sum_{\tau=t_{i,0}+1}^t \sum_{i \in \{k,a,t\}} I_{k,i,t_{i,0}} I_{i,\tau-1} \cdot X'_{i,\tau-1} \hat{\Gamma}_{k,k}, \quad (1.62)$$

where in the second equation $\sum_{i \in \{k,a,t\}}$ is a shorthand for summing over all workers of age a in year t whose initial occupation was k , and $N_{k,a,t}$ is the total number of such workers. As before, $t_{i,0} \leq t$ denotes the year when worker i joins the labor market. In Equation (1.61) we then weigh by the relative cell sizes $P_t(a, k)$ to obtain the average skill accumulation in the respective wage percentile.⁶³

Figure 1.56: Skills accumulated during working life by percentile of wage distribution



Notes: Estimates for average skill accumulation obtained in the population by year and percentile in the wage distribution following Equations 1.61, 1.62 and 1.63.

The black solid line in Figure 1.56 depicts this average skill accumulation across the percentiles of the 2010 wage distribution. We can see that skill accumulation’s

⁶³An index for the specific wage percentile is omitted, since this is always conditioned on anyway.

contribution to log wages is substantially higher at the median than at the bottom of the distribution, and much higher at the top of the distribution. The gray solid line depicts the corresponding skill accumulation for the year 1985, which is substantially flatter in its lower half but comparably steep as the 2010 skill accumulation between the 50th and the 85th percentile. The difference between the two lines is the effect of the accumulation (i.e., Equation (1.57) compared to the scenario with only initial wages changing in Equation (1.56)).

Using Equations (1.61)–(1.62), we now decompose the difference between the 2010 and 1985 skill accumulation effect into its parts with a particular focus on the lower half. One obvious component of this is supposed to be the changing occupation structure. Using Bayes' law, we compute the accumulation that would have prevailed if (conditionally) the age structure and within-cell accumulation changed over time but the (initial) occupation structure had stayed the same as in 1985:

$$avg_{2010}(occup = 1985) = \sum_{k=1}^{120} \sum_{a=25}^{54} P_{1985}(k) \cdot P_{2010}(a|k) \cdot \overline{acc}_{k,a,2010} \quad (1.63)$$

The pink dashed line depicts this skill accumulation, showing that it actually does not explain any of the increase in lower-half inequality. The reason for this rather surprising result is that the occupation structure actually did not shift decidedly toward higher-accumulation occupations at the median. Figure 1.57 depicts the corresponding graph in the scenario at hand, i.e., with workers' initial occupations in the distribution of wages only due to entry and skill accumulation. We see that there are the same share of high-accumulation occupations at the median in 1985 as in 2010.⁶⁴

The next potential component of the skill accumulation effect is the shifting age structure in the different parts of the wage distribution. That is, we change the unconditional age distribution at each percentile to its 1985 value but hold the accumulation and the conditional occupation structure at their 2010 values:

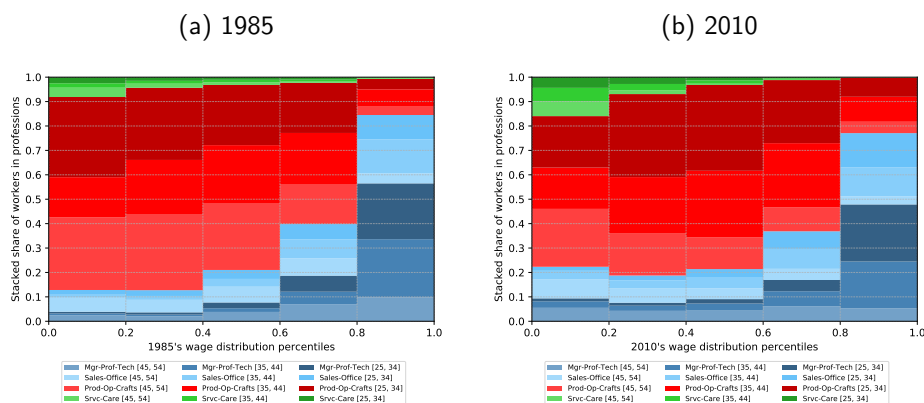
$$avg_{2010}(age = 1985) = \sum_{k=1}^{120} \sum_{a=25}^{54} P_{1985}(a) \cdot P_{2010}(k|a) \cdot \overline{acc}_{k,a,2010} \quad (1.64)$$

The yellow dashed line in Figure 1.56 shows that this strikingly explains more than half of the accumulation effect. In particular, the many experienced workers at the 2010 median, already mentioned in the main text, have accumulated a lot of skills over their careers. Figure 1.57 illustrates this even more clearly in the scenario at hand: the share of 45–54 (but also 35–44) year old Prod-Op-Crafts workers at the median is very large in 2010 and much higher than in 1985. Hence, if we down-weight the share of these workers to 1985 and hold everything else constant, as we do in Equation (1.64), skill accumulation at the median is substantially lower. Therefore, the large baby boomer birth cohorts, who still started their careers in Prod-Op-Crafts occupations,⁶⁵

⁶⁴Admittedly, this may also partially be a reflection of data limitations here, since we need to approximate initial occupations in 1985 by their 1975 occupations for workers who joined the labor market before that (see also discussion further below). That is, some (middle-aged or older) 'blue' Mgr-Prof-Tech or 'Sales-Office' workers in 1985 may have started in the 'red' or 'green' occupations before 1975. This would also explain the higher share of 'blue' initial occupations at the top quintile of the Figure 1.57 wage distribution in 1985 than in 2010 and the concurrent stronger 85th percentile increase in the pink 1985 occupation structure series of Figure 1.56.

⁶⁵The baby boom in Germany started later than in the U.S., with cohorts comprising birth years 1955–69, i.e., 41–55 year olds in 2010.

Figure 1.57: Shares in the wage distribution by quintile



Notes: Each rectangle is proportional to the share of workers in the respective occupation \times age bin. Wage quintiles were computed with wages computed as in Equation (1.57), i.e. wage growth only occurring because of skill accumulation.

are substantially raising median wages at the end of our analysis period. Bottom wage earners in 2010 are in contrast much younger, and therefore skill accumulation due to demographic change raises lower-half wage inequality in this point in time.

The last factor that may have changed between 1985 and 2010 is the skill accumulation within worker-cells, which we analyze by computing:

$$avg_{2010}(accum = 1985) = \sum_{k=1}^{120} \sum_{a=25}^{54} P_{2010}(a, k) \cdot \overline{acc}_{k,a,1985}. \quad (1.65)$$

This counterfactual represents a specific version of changing worker employment biographies: since we condition on a fixed initial occupation and age, it can only differ when workers on average have more or less gap years of not working or different ages at labor market entry in 1985 than in 2010. In Equation (1.62) above, the former corresponds to differences in the number of $I_{i,\tau-1} = 0$ instances (e.g., due to unemployment) and the latter to differences in the total number of years $t - t_{i,0}$ over which skill accumulation is summed for a given age.

The differences in labor market entry ages are hard to measure in our data because we have to impute them for workers who joined the labor market before the beginning of our sample in 1975, that is, who already appear in the 1975 data aged older than 25. We do this by computing, for every occupation, the average entry age across all years from 1976 onward and assign this as the entry age together with their 1975 occupation to every worker who appears in the sample older than 25 in that year. We then impute the wage at labor market entry by subtracting the respective skill accumulation coefficients back to that entry age. This imputation, which affects our computed 1985 (but not 2010) skill accumulation,⁶⁶ may generate a bias. We assess this possibility by pretending we only observed 2010 workers' labor market histories after 2001 (i.e., as in 1985, everything longer than nine years ago is unobserved) and then conducting the

⁶⁶Given that workers exit the sample at age 54, after the year 2003 nobody is imputed anymore.

same imputation.⁶⁷

Since the red dashed line is below the solid black line, Figure 1.56 shows that the imputation does indeed underestimate the skill accumulation in 2010, especially at and above the median. Therefore, part of our measured skill accumulation effect on lower-half inequality between 2010 and 1985 may be due to this data issue. Nonetheless, if we apply our calculation (1.64) on top of the imputed labor market biographies in 2010, age structure differences between the beginning and the end of the analysis period still have a substantial effect on 50–15 inequality. In fact, the combination of imputation and age structure (green dotted line) is almost exactly the same as the actual 1985 skill accumulation in the lower half of the wage distribution.

Finally, we examine the effect of potentially more intermittent prior labor market attachment at either end of the analysis period. That is, similar to Section 1.F.1 we fill up gaps in employment biographies (e.g., due to unemployment) during the nine years leading up to 1985 and 2010, respectively. The dotted blue line shows the results and that actually this has no discernible effect on skill accumulation in addition to the imputation and the age structure effect. In unreported analyses, we found that labor market biographies were indeed somewhat more intermittent leading up to 2010 than 1985, but this only affected the very bottom of the wage distribution (below the 5th percentile) and turns out quantitatively unimportant here.

To summarize the overall result, had we plotted

$$avg_{2010}(age = 1985, accum = 1985) = \sum_{k=1}^{120} \sum_{a=25}^{54} P_{1985}(a) \cdot P_{2010}(k|a) \cdot \overline{acc}_{k,a,1985} \quad (1.66)$$

into the Figure 1.56, it would have almost exactly overlapped with the green dotted line. This implies that the workforce's changing age structure was the main driver of rising lower-half inequality. There is in fact no role for the occupation distribution at labor market entry *conditional on* the age structure (i.e., replacing $P_{2010}(k|a)$ with its 1985 value does not matter for the lower half). Replacing $\overline{acc}_{k,a,1985}$ with a version of $\overline{acc}_{k,a,2010}$ in which initial occupations and wages are imputed as in 1985 also gives the same result (which is the green dotted line). On top of that, filling labor market biographies during the preceding nine years also does not matter (blue dotted line). Therefore, some of the lower-half inequality effect may or may not be attributed to changing initial wages instead of skill accumulation, but the economically and quantitatively important part is accounted for by aging of the workforce.

⁶⁷Conversely, we have also not imputed at all and taken the first observed wage as the initial one, which actually decreased 50–15 differences in 1985 and thus modestly raised the increase of lower-half inequality in the $\hat{w}_{i,t}^e$ scenario (i.e., lowered the part that changing skill accumulation accounts for).

References

- Abowd, John M, Francis Kramarz, and David N Margolis (1999). "High Wage Workers and High Wage Firms". In: *Econometrica* 67.2, pages 251–333.
- Acemoglu, Daron and David H. Autor (2011). "Skills, Tasks and Technologies: Implications for Employment and Earnings". In: *Handbook of Labor Economics, vol. 4 Part B*. Edited by David Card and Orley Ashenfelter. Elsevier. Chapter 12, pages 1043–1171.
- Altonji, Joseph G and Charles R Pierret (2001). "Employer Learning and Statistical Discrimination". In: *Quarterly Journal of Economics* 116.1, pages 313–350.
- Anderson, T. W. and Cheng Hsiao (1982). "Formulation and Estimation of Dynamic Models Using Panel Data". In: *Journal of Econometrics* 18.1, pages 47–82.
- Antonczyk, Dirk, Bernd Fitzenberger, and Ute Leuschner (2009). "Can a Task-Based Approach Explain the Recent Changes in the German Wage Structure?" In: *Jahrbücher für Nationalökonomie und Statistik* 229.2-3, pages 214–238.
- Arellano, Manuel and Stephen Bond (1991). "Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations". In: *Review of Economic Studies* 58.2, pages 277–297.
- Artaç, Erhan and John McLaren (2015). "Trade Policy and Wage Inequality: A Structural Analysis with Occupational and Sectoral Mobility". In: *Journal of International Economics* 97.2, pages 278–294.
- Autor, David (2019). "Work of the Past, Work of the Future". In: *Richard T. Ely Lecture given at Annual Meeting of American Economics Association, Atlanta, GA, January*. Volume 4.
- Autor, David H. and David Dorn (2013). "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market". In: *American Economic Review* 103.5, pages 1553–1597.
- Autor, David H., David Dorn, and Gordon H. Hanson (2013). "The China syndrome: Local labor market effects of import competition in the United States". In: *American Economic Review* 103.6, pages 2121–2168.
- Autor, David H., Lawrence F. Katz, and Melissa S. Kearney (2008). "Trends In U.S. Wage Inequality: Revising the Revisionists". In: *Review of Economics and Statistics* 90.2, pages 300–323.
- Autor, David H., Frank Levy, and Richard J. Murnane (2003). "The Skill Content of Recent Technological Change: An Empirical Exploration". In: *Quarterly Journal of Economics* 118.4, pages 1279–1333.
- Autor, David, David Dorn, Gordon Hanson, and Kaveh Majlesi (2016). "Importing Political Polarization? The Electoral Consequences of Rising Trade Exposure". Working Paper.

- Bárány, Zsófia and Christian Siegel (2018). "Job Polarization and Structural Change". In: *American Economic Journal: Macroeconomics* 10.1, pages 59–89.
- Battisti, Michele, Christian Dustmann, and Uta Schönberg (2017). "Technological and Organizational Change and the Careers of Workers". Working Paper.
- Baumgarten, Daniel and Sabine Lehwald (2019). "Trade Exposure and the Decline in Collective Bargaining: Evidence from Germany". Working Paper.
- Black, Sandra E. and Alexandra Spitz-Oener (2010). "Explaining Women's Success: Technological Change and the Skill Content of Women's Work". In: *Review of Economics and Statistics* 92.1, pages 187–194.
- Blinder, Alan S and Alan B Krueger (2013). "Alternative Measures of Offshorability: A Survey Approach". In: *Journal of Labor Economics* 31.2 Part 2, S97–S128.
- Böhm, Michael (2019). "The Price of Polarization: Estimating Task Prices Under Routine-Biased Technical Change". Working Paper.
- Card, David, Jörg Heining, and Patrick Kline (2013). "Workplace Heterogeneity and the Rise of West German Wage Inequality". In: *Quarterly Journal of Economics* 128.3, pages 967–1015.
- Carneiro, Pedro and Sokbae Lee (2011). "Trends in Quality-Adjusted Skill Premia in the United States, 1960-2000". In: *American Economic Review* 101.6, pages 2309–2049.
- Case, Anne and Angus Deaton (2015). "Rising Morbidity and Mortality in Midlife among White Non-Hispanic Americans in the 21st Century". In: *Proceedings of the National Academy of Sciences* 112.49, pages 15078–15083.
- Cavaglia, Chiara and Ben Etheridge (2017). "Job Polarization, Task Prices and the Distribution of Task Returns". Working Paper.
- Chetty, Raj, John N. Friedman, Nathaniel Hendren, Maggie R. Jones, and Sonya R. Porter (2018). "The Opportunity Atlas: Mapping the Childhood Roots of Social Mobility".
- Chetty, Raj, John N. Friedman, Nathaniel Hilger, Emmanuel Saez, Diane Whitmore Schanzenbach, and Danny Yagan (2011). "How Does Your Kindergarten Classroom Affect Your Earnings? Evidence from Project STAR". In: *Quarterly Journal of Economics* 126.4, pages 1593–1660.
- Combes, Pierre-Philippe, Gilles Duranton, and Laurent Gobillon (2008). "Spatial Wage Disparities: Sorting Matters!" In: *Journal of Urban Economics* 63.2, pages 723–742.
- Cornelissen, Thomas, Christian Dustmann, Anna Raute, and Uta Schönberg (2018). "Who Benefits from Universal Child Care? Estimating Marginal Returns to Early Child Care Attendance". In: *Journal of Political Economy* 126.6, pages 2356–2409.

- Cortes, Guido Matias (2016). "Where Have the Middle-Wage Workers Gone? A Study of Polarization Using Panel Data". In: *Journal of Labor Economics* 34.1, pages 63–105.
- Cortes, Guido Matias and Giovanni Gallipoli (2017). "The Costs of Occupational Mobility: An Aggregate Analysis". In: *Journal of the European Economic Association*, pages 275–315.
- Dauth, Wolfgang, Sebastian Findeisen, Jens Südekum, and Nicole Woessner (2017). "German Robots - The Impact of Industrial Robots on Workers". Working Paper.
- Deming, David J. (2017). "The Growing Importance of Social Skills in the Labor Market". In: *Quarterly Journal of Economics* 132.4, pages 1593–1640.
- DiNardo, John, Nicole M. Fortin, and Thomas Lemieux (1996). "Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semiparametric Approach". In: *Econometrica* 64.5, pages 1001–1044.
- Dix-Carneiro, Rafael (2014). "Trade Liberalization and Labor Market Dynamics". In: *Econometrica*, pages 825–885.
- Dustmann, Christian, Bernd Fitzenberger, Uta Schönberg, and Alexandra Spitz-Oener (2014). "From Sick Man of Europe to Economic Superstar: Germany's Resurgent Economy". In: *Journal of Economic Perspectives* 28.1, pages 167–188.
- Dustmann, Christian, Johannes Ludsteck, and Uta Schönberg (2009). "Revisiting the German wage structure". In: *Quarterly Journal of Economics* 124.2, pages 843–881.
- Eichhorst, Werner and Verena Tobsch (2013). "Has Atypical Work Become Typical in Germany?" IZA Discussion Paper No. 7609.
- Fetzer, Thiemo (forthcoming). "Did Austerity Cause Brexit?" In: *American Economic Review*.
- Firpo, Sergio, Nicole M. Fortin, and Thomas Lemieux (2013). "Occupational Tasks and Changes in the Wage Structure".
- Fitzenberger, Bernd (1999). *Wages and Employment Across Skill Groups: An Analysis for West Germany*. Technical report. Heidelberg: Physica/Springer.
- Fitzenberger, Bernd, Anderonke Osikominu, and Robert Völter (2006). "Imputation Rules to Improve the Education Variable in the IAB Employment Sub-sample". In: *Schmollers Jahrbuch (Journal of the Applied Social Sciences)* 126, pages 405–436.
- Gabaix, Xavier and Augustin Landier (2008). "Why has CEO Pay Increased So Much?" In: *Quarterly Journal of Economics* 123.1, pages 49–100.
- Ganzer, Andreas, Alexandra Schmucker, Philipp vom Berge, and Anja Wurdack (2017). *Sample of Integrated Labour Market Biographies Regional File 1975-2014 (SIAB-R 7514)*. Technical report 1. FDZ-Datenreport.
- Gathmann, Christina and Uta Schönberg (2010). "How General is Human Capital? A Task-Based Approach". In: *Journal of Labor Economics* 28.1, pages 1–49.

- Gaudecker, Hans-Martin von (2014). "Templates for Reproducible Research Projects in Economics". URL: <https://github.com/hmgaudecker/econ-project-templates>.
- Gibbons, Robert, Lawrence F. Katz, Thomas Lemieux, and Daniel Parent (2005). "Comparative Advantage, Learning, and Sectoral Wage Determination". In: *Journal of Labor Economics* 23.4, pages 681–724.
- Gibbons, Robert and Michael Waldman (1999). "A Theory of Wage and Promotion Dynamics Inside Firms". In: *Quarterly Journal of Economics* 114.4, pages 1321–1358.
- Glitz, Albrecht and Daniel Wissmann (2018). "Skill Premiums and the Supply of Young Workers in Germany". Working Paper.
- Goos, Maarten and Alan Manning (January 2007). "Lousy and Lovely Jobs: The Rising Polarization of Work in Britain". In: *Review of Economics and Statistics* 89.1, pages 118–133.
- Goos, Marten, Alan Manning, and Anna Salomons (2014). "Explaining Job Polarization: Routine-Biased Technological Change and Offshoring". In: *American Economic Review* 104.8, pages 2509–2526.
- Gottschalk, Peter, David A. Green, and Benjamin M. Sand (2015). "Taking Selection to Task: Bounds on Trends in Occupational Task Prices for the U.S., 1984-2013". Working Paper.
- Green, David A and Benjamin M Sand (2015). "Has the Canadian labour market polarized?" In: *Canadian Journal of Economics* 48.2, pages 612–646.
- Groes, Fane, Philipp Kircher, and Iouri Manovskii (2014). "The U-shapes of occupational mobility". In: *Review of Economic Studies* 82.2, pages 659–692.
- Heckman, James J. and Bo E. Honoré (1990). "The Empirical Content of the Roy Model". In: *Econometrica* 58, pages 1121–1149.
- Heckman, James J., Lance Lochner, and Christopher Taber (1998). "Explaining Rising Wage Inequality: Explorations with a Dynamic General Equilibrium Model of Labor Earnings with Heterogeneous Agents". In: *Review of Economic Dynamics* 1, pages 1–58.
- Heckman, James J. and Guilherme Sedlacek (1985). "Heterogeneity, Aggregation, and Market Wage Functions: An Empirical Model of Self-Selection in the Labor Market". In: *Journal of Political Economy* 93.6, pages 1077–1125.
- Hsieh, Chang-Tai, Erik Hurst, Charles I. Jones, and Peter J. Klenow (2019). "The Allocation of Talent and U.S. Economic Growth". In: *Econometrica* 87.5, pages 1439–1474.
- Jäger, Simon and Jörg Heining (2019). "How Substitutable are Workers? Evidence from Worker Deaths". Working Paper.
- Jardim, Ekaterina, Mark C Long, Robert Plotnick, Emma van Inwegen, Jacob Vigdor, and Hilary Wething (2017). "Minimum Wage Increases, Wages, and Low-Wage Employment: Evidence from Seattle". Working Paper.

- Lagakos, David, Benjamin Moll, Tommaso Porzio, Nancy Qian, and Todd Schoellman (2018). "Life-Cycle Wage Growth Across Countries". In: *Journal of Political Economy*, pages 797–849.
- Lazear, Edward P (1981). "Agency, Earnings Profiles, Productivity, and Hours Restrictions". In: *American Economic Review* 71.4, pages 606–620.
- Lee, Donghoon and Kenneth I. Wolpin (2006). "Intersectoral Labor Mobility and the Growth of the Service Sector". In: *Econometrica* 74.1, pages 1–46.
- Liu, Kai, Kjell G. Salvanes, and Erik Sørensen (2016). "Good Skills in Bad Times: Cyclical Skill Mismatch and the Long-Term Effects of Graduating in a Recession". In: *European Economic Review* 84, pages 3–17.
- Mazzolari, Francesca and Giuseppe Ragusa (2013). "Spillovers from high-skill consumption to low-skill labor markets". In: *Review of Economics and Statistics* 95.1, pages 74–86.
- McLaughlin, Kenneth J and Mark Bilts (2001). "Interindustry Mobility and the Cyclical Upgrading of Labor". In: *Journal of Labor Economics* 19.1, pages 94–135.
- Milgrom, Paul and Ilya Segal (2002). "Envelope Theorems for Arbitrary Choice Sets". In: *Econometrica* 70.2, pages 583–601.
- Mishel, Lawrence, Heidi Shierholz, and John Schmitt (2013). "Don't Blame the Robots: Assessing the Job Polarization Explanation of Growing Wage Inequality". Working Paper.
- Naticchioni, Paolo, Giuseppe Ragusa, and Riccardo Massari (2014). "Unconditional and Conditional Wage Polarization in Europe". Working Paper.
- Newey, Whitney K. and Frank Windmeijer (2009). "Generalized Method of Moments With Many Weak Moment Conditions". In: *Econometrica* 77.3, pages 687–719.
- Papageorgiou, Theodore (2014). "Learning Your Comparative Advantages". In: *Review of Economic Studies* 81, pages 1263–1295.
- Roy, A. D. (1951). "Some Thoughts on the Distribution of Earnings". In: *Oxford Economic Papers* 3.2, pages 135–146.
- Solon, Gary, Robert Barsky, and Jonathan A Parker (1994). "Measuring the Cyclicalities of Real Wages: How Important is Composition Bias?" In: *Quarterly Journal of Economics* 109.1, pages 1–25.
- Spitz-Oener, Alexandra (2006). "Technical Change, Job Tasks, and Rising Educational Demands: Looking outside the Wage Structure". In: *Journal of Labor Economics* 24.2, pages 235–270.
- Taber, Chris and Nicolas Roys (2019). "Skill Prices, Occupations and Changes in the Wage Structure for Low Skilled Men". Working Paper.
- Yamaguchi, Shintaro (2018). "Changes in Returns to Task-Specific Skills and Gender Wage Gap". In: *Journal of Human Resources* 53.1, pages 32–70.
- Young, Alwyn (2014). "Structural Transformation, the Mismeasurement of Productivity Growth, and the Cost Disease of Services". In: *American Economic Review* 104.11, pages 3635–3667.

Chapter 2

Changing Returns to Occupational Skill and Women's Wages

2.1 Introduction

This paper investigates the importance of changing returns to occupational skill and declining occupational segregation for the reduction in wage inequality between men and women. During the last three decades, wage inequality has increased in almost all developed countries (Acemoglu and Autor, 2011). This increase has taken place within as well as between observable groups defined by variables such as education, occupation, region or industry. The major exception to the rule – some authors even argue, paradoxical exception to the rule (Blau and Kahn, 1997; Card and DiNardo, 2002) – has been the convergence of male and female wages over time (Olivetti and Petrongolo, 2016; Blau and Kahn, 2017).

Many possible explanations for the declining gender wage gap exist and have been raised in the literature such as rising female labor market experience (Olivetti, 2006), positive selection of women into the labor force (Mulligan and Rubinstein, 2008), gender biased shifts in labor demand (Heathcote et al., 2010), structural change (Ngai and Petrongolo, 2017), or declining discrimination (Hsieh et al., 2019). Apart from that, another very popular explanation is concerned with returns to skill having changed in a women benefiting way, despite these changing skill returns might have raised overall inequality (Altonji and Blank, 1999).

The idea for this goes back to Goldin (1990) who hypothesized that women have a comparative advantage in tasks that require “brains” rather than “brawns”. If changes in labor demand (for instance, due to technical change) induced an increase in the price paid for brain skills, then wage inequality *between* men and women should decline; although inequality *within* male and female samples is

likely to rise.¹

In line with this demand side driven view, this paper investigates the importance of changes in occupational demand and supply for the declining gender wage gap between men and women. Using German administrative panel data with consistent and accurate information on workers' detailed occupations and their wages, I find that almost all of the decline in the gender wage gap between 1985 and 2010 can be explained by a convergence in occupational choices between men and women as well as changing occupational wages.

The reason for this is that men and women have worked (Weisskoff, 1972) and still do work (Cortes and Pan, 2018) in very different occupations. This segregation leads to two potential sources of convergence between men's and women's wages. These sources are commonly referred to as the "wage structure effect" and the "composition effect" in the literature (Blau and Kahn, 2017): a change in the wage structure benefits women's relative wages if wages are increasing in occupations which make up a great share of female employment. In turn, a change in employment composition reduces the wage gap if the relative number of women in high wage occupations increases.

So far, these two hypothesized effects have been analyzed in isolation from each other. However, recent research suggests that changes in occupational employment growth might have a direct, attenuating effect on occupational wage growth (and thereby the wage structure effect). The reason for this is that workers self select into occupations depending on skill prices (Cortes, 2016; Böhm et al., 2019; Hsieh et al., 2019). Hence, contingent on the magnitude of workers' selection response, the wage structure effect might not reflect the causal influence of changing skill prices on changes in gender wage inequality. For instance, the selection response in the model of Hsieh et al. (2019) exactly offsets a change in skill prices so that wage and employment growth are completely uncorrelated. US data are consistent with this theoretical prediction. The same finding holds for Germany (Böhm et al., 2019). When ignoring this, the result might be an overestimation of the composition effect and an underestimation of the wage structure effect. This is because wages do not move as much as predicted by changing skill prices alone, that is, in absence of worker self selection.

Nevertheless, to identify the influence of the wage structure and composition effect, most existing work has typically relied on simple regression models for estimating occupational wage premia in a base period and then asked:² how would the wage gap have changed if only employment but not wages had changed? In turn, estimates for the wage structure effect are derived from holding employment shares constant and only letting average wages change over time. For instance, following the Juhn et al. (1993) approach, Bacolod and Blum (2010)

¹See Welch (2000) for a formalization of this argument.

²This approach goes back to a larger literature on estimating changes in wage differentials (see Smith and Welch, 1989; Juhn et al., 1991, 1993).

find that rising wages in cognitive and social occupations can explain 20% of the decline in the US gender wage gap. The bulk of the remaining convergence is explained by the changing composition of female work. Weinberg (2000) and Borghans et al. (2014) confirm this finding for the US and complement it with evidence for the UK and Germany. Black and Spitz-Oener (2010) employ data on tasks which workers execute on the job and estimate that about half of the declining gap can be traced back to an increase in the non-routine component of work.

To repeat the argument from above: what connects these approaches in identifying changes of skill prices from time series variation in wages is the implicit assumption that the quality of the average worker within an occupation must not change in response to shifting skill prices.³ In a first step, I follow this approach and decompose the change in the gender wage gap into a wage structure effect, a composition effect and a residual. Overall, I find that the gender wage gap within a sample of full time working 25–54 year old workers declined from being as high as 44 log points in 1985 to 36 log points in 2010. Roughly 35% of that decline are explained by changes in average occupational wages: women have worked and still do work in occupations in which average wages increased relatively more than in occupations in which many men work (such as manufacturing). In turn, nearly 65% of the decline can be attributed to changing employment patterns: women moved into high wage, professional and managerial occupations previously dominated by men (Ngai and Petrongolo, 2017). In fact, cohort analyses show that the proportion of women at labor market entry in managerial occupations increased dramatically over time, rising from 40% for women born between 1945–1955 to 68% for those born after 1965.

In the next step, I then move beyond equating changes in skill prices with changes in average occupational wages. I take (quality adjusted) skill price change estimates from earlier work in Böhm et al. (2019). They identify the part of average wage growth which is due to changing skill prices by employing long term panel data which allow them to hold constant occupation specific changes in worker quality. Using these skill price estimates, I decompose the wage structure effect further into a price effect and a skill selection effect. Importantly, the skill price estimates come from a sample of prime age men as the underlying (static) occupational choice model is less likely to be misspecified for men because forward looking behavior might play a less prominent role for men compared to women (Adda et al., 2017). Hence, I implicitly assume that men and women do not work in segmented labor markets but instead are paid the same price for a unit of skilled labor as it would be the case in a competitive

³Contrary to relying on time series variation in occupational wages, Beaudry and Lewis (2014) exploit variation in wage gaps together with variation in PC adoption across local labor markets. They find that the gender wage gap decreased more in areas with higher PC adoption. They interpret this result as being driven by higher increases in cognitive skill prices paid within areas that experienced more abundant PC adoption.

labor market without discrimination.⁴

The exercise shows that the raw wage structure effect indeed underestimates the contribution of changing skill prices on the convergence of male and female wages. In fact, the results suggest that changing skill prices contributed roughly 13 percentage points more to the decline in the gender wage gap than changes in average wages alone. The main reason for that is a large increase in female labor force participation during the late 1980s. By definition, this rise brought a lot of inexperienced, low wage, low skill women into work. The increasing female participation rate initially counteracted skill price movements which were essentially beneficial for women's wages. Further, the results suggest that women not only benefited from rising prices paid for tasks women perform more often than men (such as service and care), but they also benefited from declining prices paid for tasks men are engaged in more often than women (mostly manufacturing occupations). Interestingly, a similar result was found by Yamaguchi (2018) who used an alternative approach to estimate the influence of skill prices on the declining gender wage gap in US panel data (also in terms of timing).

In the last step, I investigate the importance of changing returns to occupational skill and changing employment patterns for the proportion of women at different parts of the combined male and female wage distribution. The observed proportion of women at the 85th percentile of the wage distribution increased from 12 percentage points in 1985 to 20 points in 2010. In contrast, the proportion of women at the 15th percentile was 61 percentage points in 1985. This share decreased to 50 points over time. Hence, gender wage inequality declined in both the upper and lower part of the wage distribution.

Using a reweighting approach by following DiNardo et al. (1996), I estimate how large the share of women by percentile would have been if, *ceteris paribus*, skill prices had not changed or the distribution of occupational employment had not changed over time (or both). In summary, if skill prices and employment had not changed over the years, the share of women at the 85th percentile would have increased by 50% less. Similarly, the share of women at the 15th percentile would have fallen by 55% less. Therefore, changes in prices and occupational employment induced a decline in gender wage inequality in both the upper and lower part of the wage distribution.

The remainder of this chapter is structured as follows. Section 2.2 presents the data and evidence on the declining wage gap. Section 2.3 decomposes the wage gap into wage structure, skill price and composition effects. Section 2.4 concludes.

⁴See Hsieh et al. (2019) for a setup which explicitly allows for discrimination having changed over time within a Roy model framework.

2.2 Male and Female Wage and Employment Patterns

2.2.1 Data

Data comes from the Sample of Integrated Labor Market Biographies (SIAB) Scientific Use File provided by the IAB Institute at the German Federal Employment Agency. The SIAB is a 2% random sample of administrative social security records from 1975 to 2014. It is representative of 80% of the German workforce and includes employees covered by social security, marginal part-time employment, benefit receipts, officially registered as job-seeking or participating in programs of active labor market policy. It excludes the self-employed, civil servants, and individuals performing military service. Most notably, it contains individuals' full employment histories including detailed data for workers' occupations along with socio-demographics such as age, gender, or the level of education. The contained wage measure corresponds to the daily gross wage.

I prepare the data along the lines of Böhm et al. (2019). I transfer the spell structure into a yearly panel by deleting all spells except for the longest spell within a year. I impute wages above the social security limit, and restrict the main sample to 25 to 54 year old Germans working full-time in former West Germany between 1975 and 2010.

The SIAB Scientific Use File contains information on 120 different occupations on which the decompositions are based. For some parts of the analyses, I categorize these detailed occupations into 10 broader groups broadly following Acemoglu and Autor (2011): managers, professionals, technicians, sales, office, production, operators, craftsmen, service and care occupations with managers, professionals, and technicians being high wage, analytical occupations. In contrast service and care occupations represent low wage, manual occupations whereas production, operators and craftsmen consist of middle wage, routine and manual occupations. Sales and office lie in between high and middle wage occupations. See Böhm et al. (2019) for the mapping.

In robustness checks, I relax the data restrictions to also include part-time workers. Unfortunately, the data do not contain any information on the exact number of hours so that I cannot construct hourly wages. Instead, the wage of a part-time worker will also refer to the daily wage and, hence, be lower than a full time wage because of differences in working hours. Although most of the existing literature about changes in the gender wage gap also restricts the analyses to full time working men and women (e.g., Bacolod and Blum, 2010; Black and Spitz-Oener, 2010; Blau and Kahn, 2017; Yamaguchi, 2018), I view the inclusion of part time workers as a crucial validity check because, contrary to the US, working part-time has been very prevalent for women in Germany and has even increased over time. In fact, data from the OECD (2019) show that the share of women aged 25–54 in part-time work increased from 35% in 1990 to 39% in 2010. In contrast, the proportion of women working part-time

in the US decreased from 15% to 13% over time.⁵

2.2.2 Trends in Gender Wage Inequality and Employment Gaps

Germany has one of the highest and most persistent gender wage gaps among developed countries (Olivetti and Petrongolo, 2016). Panel A in Table 2.1 shows the level and evolution of the gap for full time-working 25–54 year old men and women over time in the SIAB data. The difference in average log wages between the sexes was 44.24 log points in 1985. This gap declined by 8.42 log points until 2010 with most of the convergence in wages taking place until 1993. This timing is similar to the US (e.g., Figure 4, p. 414 in Olivetti and Petrongolo, 2016).

Table 2.1: Gender gaps in wages and occupational employment

	Level	Difference			
	1985	1993 - 1985	2001 - 1993	2010 - 2001	2010 - 1985
<i>Panel A</i>					
Observed wage gap	44.24	-4.66	-2.41	-1.34	-8.42
Pred. gap, occupation × year	20.45	-3.73	-3.13	-1.50	-8.37
Pred. gap, occupation	19.09	-2.37	-2.30	-1.37	-6.05
<i>Panel B</i>					
Duncan index	61.05	-1.51	-3.08	-1.81	-6.40
Occupation mix effect		-1.06	-1.46	-0.76	-3.28
<i>Panel C</i>					
Proportion of women					
Mgr	17.23	6.56	4.86	2.21	13.62
Prof	20.88	5.47	1.48	2.97	9.91
Tech	10.45	4.40	0.08	0.40	4.88
Sales	43.41	4.59	-1.10	-1.28	2.21
Office	64.99	3.15	-3.31	-2.75	-2.91
Prod	16.79	-0.27	-1.73	-3.28	-5.29
Op	9.16	1.61	1.03	1.19	3.84
Crafts	10.48	-0.40	-0.98	-0.98	-2.35
Srvc	58.12	-0.79	-5.19	-3.39	-9.37
Care	85.20	1.14	-1.96	-0.20	-1.02

Notes: Panel A shows log wages of men minus log wages of women for three different wage measures: observed wages, predicted wages from a regression of wages on 120 occupation dummies interacted with year dummies as well as predictions from a regression of wages on occupation dummies only. Panel B shows the Duncan index computed as in Equation (2.1). Occupation mix effect refers to the change in the Duncan index attributable to economy wide changes in occupation sizes, see Blau, Brummund, et al. (2013) for the details. Panel C shows the proportion of women within professions: Mgr: managers; Prof: professions; Tech: technicians; Prod: production; Op: Operators; Crafts: craftsmen; Srvc: services. Classification of detailed occupations follows Acemoglu and Autor (2011), see Böhm et al. (2019) for the exact mapping. Values ×100. Full-time working 25-54 year old West German men and women.

When I replace individual wages with average occupational wages by means of the prediction from a regression of log wages on occupation times year dum-

⁵The proportion of prime age men working part-time is quite low in both countries although having slightly increased between 1990 and 2010 from 2% to 6% (Germany) and from 3% to 4% (US).

mies, the resulting wage gap amounts to 20.45 log points in 1985. Hence, compared to the overall gap, 46% of the gender wage gap in 1985 can be explained by differences in the occupational structure between men and women. This already shows the importance of occupations for gender wage inequality. What is even more noteworthy, however, is the relevance of *changes* in the occupation structure for the *change* in the gender wage gap.

The last column of Table 2.1 shows changes between 1985 and 2010. To repeat, the observed log wage gap shrunk by 8.42 log points. In turn, the decrease in the wage gap predicted by changing occupational employment as well as changing occupational wages amounts to 8.37 points coming very close to the observed decline. This implies that the overwhelming majority of the closing of the gap can be attributed to either converging occupational employment of men and women (composition effect) or to relative wage increases of occupations which are important for female employment (wage structure effect). As a first pass, to distinguish the influence of the wage structure effect from the composition effect, Table 2.1 also shows the change in the wage gap predicted by a regression of wages on occupation dummies alone. The result of that exercise is informative about the influence of changes in occupational segregation, i.e. the composition effect. The predicted decline in the wage gap amounts to 6.05 log points which is substantial compared to the overall decrease.

In turn, the difference between the two predicted declines is informative about the wage structure effect. The comparably small magnitude of $8.37 - 6.05 = 2.32$ log points might, however, underestimate the influence of changing skill prices as workers in occupations are a self selected group. Section 2.3 makes this more formal. Before that, I will provide some more information on differences in employment between the sexes as well as information on how wage and employment gaps evolve over the life cycle.

Panel B of Table 2.1 shows level and changes of a segregation index originally proposed by Duncan and Duncan (1955). The index ranges between zero and one:

$$\text{Duncan index} = \frac{1}{2} \sum_k |P_t(k|m) - P_t(k|f)| \quad (2.1)$$

$P_t(k|g)$ denotes the share of employment in occupation k at time t for gender g . This index is zero if the share of men and women in every occupation equals the share of men and women in the overall working population and hence there is no segregation. In turn, the Duncan index is one if men and women work in completely distinct occupations.

In 1985, the Duncan index amounted to 61%.⁶ The interpretation of this number is that 61% of men or women would have to switch occupations for the occupational employment distribution to be the same across sexes. Until 2010,

⁶Cortes and Pan (2018, Fig. 18.1, p.427) compute an index of roughly 56% in 1985 for the US with a continuous decline to 51% in 2009.

the index decreased by 6%. There are two possible reasons for this decrease in segregation as noted by Fuchs (1975) and Blau, Brummund, et al. (2013). First, the aggregate size of previously large, male dominated occupations might have declined over time. In fact, when holding the proportion of women within occupations constant, I find that roughly 3% out of 6% are attributable to this aggregate change in occupation structure. This is mainly because of shrinking manufacturing occupations with large shares of men. Second, the proportion of women increased in occupations which once exhibited small female proportions. Panel C in Table 2.1 shows that this primarily happened within high wage managerial and professional occupations in which the share of women increased by roughly 14 and 10 percentage points. In addition, the proportion of women in service occupations declined over time so that more than 50% of (full-time) service workers today are men compared to only 42% in 1985.

Table 2.2: Wage and employment gaps when including part-time workers

	Level	Difference			
	1985	1993 - 1985	2001 - 1993	2010 - 2001	2010 - 1985
<i>Panel A</i>					
Observed wage gap	59.09	-4.48	-3.57	0.27	-7.79
Pred. gap, occupation \times year	31.59	-3.82	-3.98	-0.38	-8.17
Pred. gap, occupation	30.41	-3.04	-3.08	-1.69	-7.81
<i>Panel B</i>					
Duncan index	63.58	-1.95	-3.19	-1.65	-6.79
Occupation mix effect		-1.56	-1.61	-1.18	-4.36
<i>Panel C</i>					
Proportion of women					
Mgr	19.22	7.21	5.59	3.75	16.55
Prof	26.63	6.04	3.96	5.07	15.07
Tech	12.58	5.13	0.95	1.16	7.24
Sales	53.26	4.67	-0.58	0.81	4.90
Office	71.78	2.87	-2.81	-1.14	-1.08
Prod	19.17	-0.36	-1.84	-2.78	-4.97
Op	15.74	1.41	0.05	1.20	2.66
Crafts	12.49	-0.80	-0.99	-0.60	-2.38
Srvc	71.78	-1.48	-3.86	-1.73	-7.07
Care	88.05	1.77	-0.93	0.39	1.22

Notes: Panel A shows log wages of men minus log wages of women for three different wage measures: observed wages, predicted wages from a regression of wages on 120 occupation dummies interacted with year dummies as well as predictions from a regression of wages on occupation dummies only. Panel B shows the Duncan index computed as in Equation (2.1). Occupation mix effect refers to the change in the Duncan index attributable to economy wide changes in occupation sizes, see Blau, Brummund, et al. (2013) for the details. Panel C shows the proportion of women within professions: Mgr: managers; Prof: professions; Tech: technicians; Prod: production; Op: Operators; Crafts: craftsmen; Srvc: services. Classification of detailed occupations follows Acemoglu and Autor (2011), see Böhm et al. (2019) for the exact mapping. Values $\times 100$. Full-time and part-time working 25-54 year old German men and women. The wage of a part-time worker is not adjusted for hours. Instead, the wage refers to a working day as is also the case for full-time workers.

As the number of part time workers has increased a lot in Germany over the last decades, Table 2.2 repeats the exercise including part-time workers. As the SIAB data do not contain any information on the exact number of hours, I

cannot adjust wages for working times. Instead, the wage of a part-time worker also refers to the daily wage. Therefore, it will also be lower than the wage of full time workers because of lower hours worked (see Section 2.2.1).

The level of the wage gap becomes substantially larger when including part-time workers. It rises from a baseline of 44 log points to 59 log points in 1985 since working part time is much more common for women. The part of the gap explained by differences in occupations becomes disproportionately larger as well, however, increasing to 32 log points compared to 20 log points in the full-time sample. Additionally, both the change in the observed gap as well as the change in the predicted gap are very similar to the decline of the wage gap when excluding part-time workers. Hence, not only for the subsample of full-time workers, changes in occupational demand and supply seem to have reduced the gender wage gap; but these changes also seem to have been important for the relative wage changes of part-time working women.

2.2.3 Life Cycle Employment and Wage Profiles

The average gender wage gap masks substantial heterogeneity between cohorts as well as over the life cycle. Figure 2.1a shows the evolution of the wage gap by age for several cohorts.⁷ Starting at low levels ranging between 10 and 26 log points at age 25, the wage gap roughly doubles for every given cohort in the sample with the highest gaps observed between ages 35–40; a point in the life cycle where average male wages are between 41 and 54 log points higher than average female wages. The shape of the life cycle profiles seems not to have changed much between the cohorts except for the intercept continuously declining. Hence, most of the decline in the wage gap is due to between cohort effects whereas most of the level of the wage gap is best explained by sources happening within cohorts.⁸

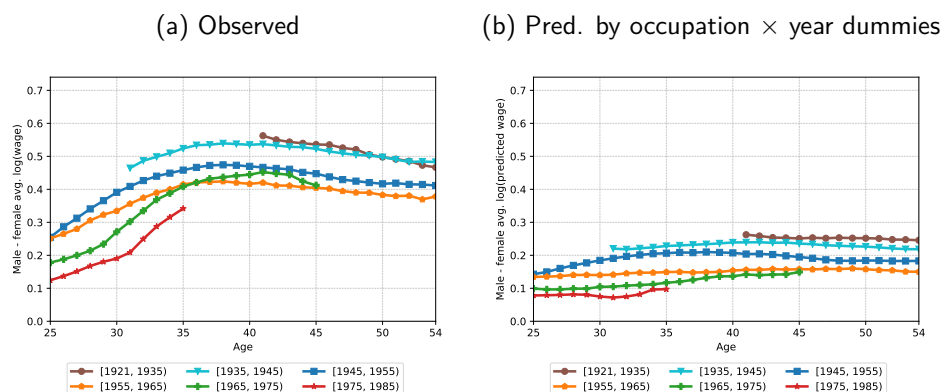
Figure 2.1b plots the predicted wage gap for different cohorts from a regression of log wages on occupation times year dummies as shown in Table 2.1. The shape of the wage gap profiles becomes rather flat. This implies that changes in the occupational structure (over time) are not very informative about the rising wage gap over the life cycle. For instance, women always might have switched to lower ranked occupations after giving birth; but this behavior was the same in 1985 as it was in 2010 (similar to the findings in Kleven et al., 2018). Nevertheless, the predicted values match the observed decline in the intercepts and therefore the overall decline of the wage gap very well.

One important reason for this has been the changing female occupation distribution – at least at the start of women’s working life cycle. Figure 2.2

⁷See Appendix Figure 2.4 for the results including part-time workers.

⁸Notice that this is fully consistent with Kleven et al. (2018, p. 1) who find that the child penalty has not changed over time despite diminishing gender inequality: “we provide a simple explanation for the persistence of gender inequality: the effects of children on the careers of women relative to men are large and have not fallen over time”.

Figure 2.1: Wage gap by cohort



Notes: Figure 2.1a shows log wages of men minus log wages of women by cohort and age. Figure 2.1b repeats this but uses the predictions from a regression of log wages on occupation times year dummies to calculate the wage gap. Full-time working 25-54 year old West German men and women.

shows the proportions of women in two high wage occupation groups, managers and professionals, for different cohorts. There is a striking rise in the share of women at the beginning of a cohort's life cycle over time. Whereas the share of women in managerial occupations was 40% for the 1945–1955 born, this proportion increased up to 68% for the cohorts born after 1965. A similar increase, though smaller, is observable for professional occupations. As cohorts age, however, the proportion of women drops significantly to around 15 – 25%. In part, this reflects decreasing female labor participation in fertile ages: the overall proportion of full-time working women shrinks from around 46% at age 25 to 30% after age 35. Nevertheless, the decrease of the proportion of women is much stronger in managerial and professional occupations compared to other occupations (for related findings see also Kleven et al., 2018). For comparison, Appendix Figure 2.5 shows the remaining broad profession groups which feature much less distinct declines.

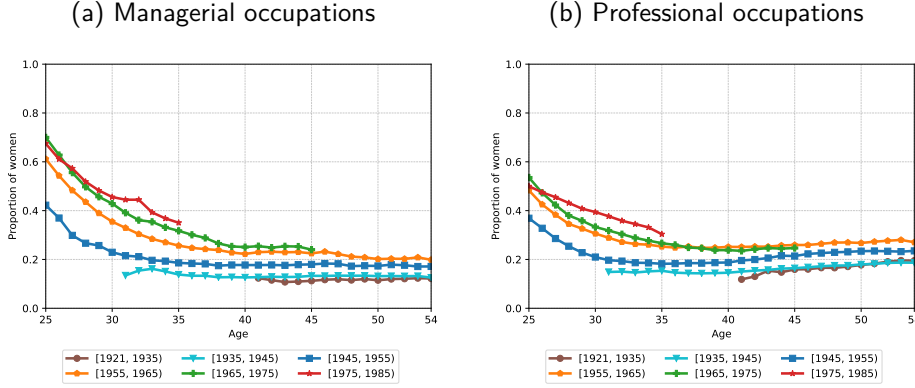
In summary, most of the decline of the wage gap has taken place between cohorts as the occupational structure of more recent cohorts is very different from the structure of older cohorts.

2.2.4 Distinguishing Changes in Average Wages from Changing Skill Prices

In the next section, I will decompose the changing wage gap into changes in the returns to skill, changes in skill selection, and changes in employment. In this section, I will highlight why it might be important to explicitly distinguish changes in average wages and changes in skill prices from each other.

Several authors starting with Goldin (1990), Welch (2000), and Weinberg (2000) have noted that advances in technology over the last few decades should

Figure 2.2: Proportion of women in professions by cohort



Notes: Classification of detailed occupations follows Acemoglu and Autor (2011), see Böhm et al. (2019) for the exact mapping. Full-time working 25-54 year old West German men and women.

have been beneficial for female labor market outcomes. For instance, Welch (2000, p. 1) argues: “In any case, if women are relatively intensive in intellectual skills and the value of such skills increase, then women’s wages will increase relative to men’s”.

To estimate the impact of changing returns to skill on the change of the gender wage gap $\Delta[\bar{w}_t^m - \bar{w}_t^f]$, with Δ denoting changes between t and $t - x$, most of the recent literature has typically decomposed changes in the gap into a part explained by the changing employment composition of men and women (or other observables), a part explained by differential changes in wage returns to occupations, and a residual part describing differences in male and female wages within occupations:

$$\Delta[\bar{w}_t^m - \bar{w}_t^f] = \underbrace{\Delta[\bar{w}_t^m - \bar{w}_t^f]}_{\text{between occupations}} + \underbrace{\Delta[\bar{e}_t^m - \bar{e}_t^f]}_{\text{within occupations}} \quad (2.2)$$

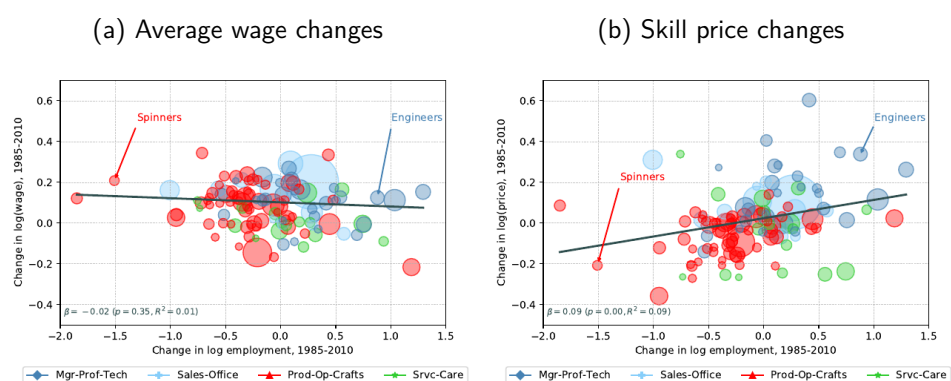
$$\begin{aligned} &= \sum_{k=1}^K \underbrace{\Delta \bar{w}_{k,t} [p_{k,t}^m - p_{k,t}^f]}_{\text{wage structure effect}} + \sum_{k=1}^K \underbrace{\bar{w}_{k,t-x} [\Delta p_{k,t}^m - \Delta p_{k,t}^f]}_{\text{composition effect}} \\ &+ \Delta[\bar{e}_t^m - \bar{e}_t^f] \end{aligned} \quad (2.3)$$

The wage structure effect holds constant occupational employment shares $p_{k,t}^g = P(k|g,t)$ of men and women over time, and so estimates how the overall gap would have changed, had only average wages $\bar{w}_{k,t}$ changed. In contrast, the composition effect holds constant average wages at a given time point to compute how the gap would have changed, had only the distribution of men and

women across occupations changed.⁹

Typically, the literature has found modest results for the wage structure effect and large composition effects. For instance, Bacolod and Blum (2010) find that only 20% percent of the narrowing US gender wage gap can be explained by changing wage returns. The opposite is true for work by Yamaguchi (2018), however, who finds that changes in the returns to motor skills were most important for the decline in the US wage gap. The main reason for the difference is that Yamaguchi (2018) allows average wages to also change because of changing skills motivated by Roy (1951): if the price paid for a task increases, workers will self select into performing more of that task; even if they are endowed with less skills in that task than in their origin occupation. Hence, wage returns might not necessarily reflect how skill prices change over time depending on workers' selection response. Therefore, estimates of the impact of changing skill prices may be heavily biased when (selection confounded) changes in average wages are used as a proxy for changes in skill prices.

Figure 2.3: Correlating employment, wage, and skill price changes



Notes: The vertical axis in Figure 2.3a shows the change in average log wages between 1985 and 2010. The vertical axis in Figure 2.3b shows the change in selection corrected skill prices between 1985 and 2010. Results are taken from Böhm et al. (2019). The horizontal axes depict the change in log employment between 1985 and 2010. One bubble represents one of the 120 detailed occupations in the SIAB SUF. The four groups show an aggregation of these detailed occupations. Classification of detailed occupations follows Acemoglu and Autor (2011), see Böhm et al. (2019) for the exact mapping. Bubble size corresponds to the number of employed men and women in an occupation averaged across years 1985 until 2010. The regression line is weighted by bubble size. Full-time working 25-54 year old West German men and women.

In line with that, Böhm et al. (2019) show that occupational wage and employment changes are essentially uncorrelated for both men and women in Germany. This raises serious doubts about the equivalence between changes in

⁹Notice that the point in time (i.e., t or $t - x$), at which wages or employment are held constant, is arbitrary in these types of decompositions. For that, in practice, I follow Neumark (1988) and use an average decomposition where I take both $t - x$ and t as base years and then average the results from the two decompositions. This does not substantially affect the magnitude of the estimates.

2.3. IMPACT OF CHANGING SKILL RETURNS ON THE GENDER WAGE GAP¹⁶⁷

skill prices and changes in average wages. Figure 2.3a shows that the coefficient from a regression of average wage changes between 1985 and 2010 on log employment changes is $\beta = -0.02$ with a p -value of 0.35. This non-correlation might be due to large selection effects taking place in the labor market which offset changing skill returns within a Roy model setup (see also Hsieh et al., 2019).

In fact, when holding constant the skill distribution over time, selection corrected skill prices $\Delta\hat{\pi}_{k,t}$ and employment are highly correlated. The regression coefficient is $\beta = 0.09$ (p -value < 0.01).¹⁰

According to the Roy framework in Böhm et al. (2019), the difference between average wage changes and skill price changes is attributed to changes in average skills $\Delta\hat{s}_{k,t}$ which counteract the changing skill prices because of worker self selection:

$$\Delta\bar{w}_{k,t} = \Delta\hat{\pi}_{k,t} + \Delta\bar{s}_{k,t} \quad (2.4)$$

Hence, it might be very important to distinguish wage changes that occur because of changing skill prices from wage changes which appear because of a changing skill selection for the estimation of the wage structure effect. For instance, such an exercise is informative about the question: did the gap close because workers became more skilled over time in occupations important for female employment; or did the gap close because of rising returns to skill within occupations important for female employment? Only the latter effect would be informative about the role of demand changes for the declining gender wage gap. The next section investigates this in more detail.

2.3 Impact of Changing Skill Returns on the Gender Wage Gap

The previous section showed that most of the decline in the German gender wage gap is explained by changes in the occupational structure between men and women. On the one hand, the proportion of women employed in high wage occupations increased substantially over time. On the other hand, changes in average wages were also beneficial for women's wages compared to men. The literature has attributed a changing occupational wage structure to shifts in labor demand moving away from male dominated occupations which require physical strength and so are less suited for women (Acemoglu and Autor, 2011). However, changes in average wages not only reflect the influence of changing skill prices; but they also reflect how the quality of workers within these occupations

¹⁰The variation in the data causing the positive relationship between employment and skill price growth comes from growing occupations having accelerating individual wage growth relative to a base period. A difference between average individual wage growth and changes in average wages partly arises because of the mechanically large number of workers entering growing occupations as well as low wages of these entrants relative to incumbents.

changes due to supply effects (Roy, 1951). This section explicitly separates these two explanatory approaches to get a more comprehensive estimate for the influence of changing skill prices on the declining gender wage gap.

2.3.1 Decomposition of the Average Gender Wage Gap

In the Roy type framework of Böhm et al. (2019), average wages either change because of changing skill prices or because of changing skills workers possess. This allows me to rewrite the wage structure effect of Equation (2.4) as:

$$\sum_{k=1}^K \underbrace{\Delta \bar{w}_{k,t} [p_{k,t}^m - p_{k,t}^f]}_{\text{wage structure effect}} = \sum_{k=1}^K \underbrace{\Delta \hat{\pi}_{k,t}^{>0} [p_{k,t}^m - p_{k,t}^f]}_{\text{growing prices}} + \sum_{k=1}^K \underbrace{\Delta \hat{\pi}_{k,t}^{\leq 0} [p_{k,t}^m - p_{k,t}^f]}_{\text{declining prices}} \quad (2.5)$$

$$+ \sum_{k=1}^K \underbrace{\Delta \bar{s}_{k,t} [p_{k,t}^m - p_{k,t}^f]}_{\text{skill changes}}$$

Table 2.3 shows the results of this decomposition. As already noted before, the between occupation gap declined by 8.37 log points between 1985 and 2010. Roughly 65% of this between effect can be explained by changes in occupational employment whereas 35% can be attributed to changes in average occupational wages.

However, the overall wage structure effect differs from the effect due to changing skill prices; and therefore differs from the causal effect of changing skill prices. In fact, decreasing prices in some occupations contributed the greatest part to the decline inducing the gender wage gap to decrease by minus three log points. The main cause for this are declining prices in crafts and production occupations. See Appendix Table 2.5 which breaks up the impact of changing prices by broad profession group. Rising prices, in turn, contributed another one log point to the decline mostly because of growing prices in sales and office occupations with large proportions of women. Hence, the causal effect of changing skill prices on the the decline in the gender wage gap was 13% larger than the raw wage structure effect.

The reason for this is that changing skills of men and women increased the gender wage gap in the aggregate thereby mitigating the wage structure effect. Improving male skills in production and crafts occupations were mostly responsible for this effect (see Appendix Table 2.5). This is in line with the strong positive selection effects in declining occupations found by Böhm et al. (2019).¹¹

¹¹Appendix Table 2.6 repeats the analysis including part-time workers. The inclusion does not change the results qualitatively but the skill change estimate rises up to an overall effect of 3.44 log points between 1985 and 2010. The reason for this is the increasing number of female part-time workers with lower wages than full time workers. Notice that the decomposition attributes these lower wages to lower skills although differences in wages between part-time and full-time workers also reflect differences in hours because the wage measure in the SIAB refers to the daily, not hourly, wage.

2.3. IMPACT OF CHANGING SKILL RETURNS ON THE GENDER WAGE GAP 169

Table 2.3: Effect of changing skill prices on gender wage inequality

	Level	Difference			
	1985	1993 - 1985	2001 - 1993	2010 - 2001	2010 - 1985
<i>Panel A</i>					
Observed gap	44.24	-4.66	-2.41	-1.34	-8.42
Pred. gap, occupation \times year	20.45	-3.73	-3.13	-1.50	-8.37
<i>Panel B</i>					
Wage structure effect		-1.44	-0.97	-0.55	-2.96
Growing prices		-1.47	0.23	0.21	-1.03
Declining prices		-3.49	0.35	0.14	-3.01
Skill changes		3.52	-1.54	-0.89	1.08
Composition effect		-2.29	-2.17	-0.95	-5.41
<i>Panel C</i>					
Within occupation gap	23.79	-0.93	0.72	0.16	-0.06
$\Delta p_{k,t}^m > \Delta p_{k,t}^f$		-1.38	-0.67	-0.75	-2.80
$\Delta p_{k,t}^m \leq \Delta p_{k,t}^f$		0.45	1.39	0.91	2.74

Notes: Panel A shows log wages of men minus log wages of women for two different wage measures: observed wages and predicted wages from a regression of wages on 120 occupation dummies interacted with year dummies, i.e. the between occupation gap in Equation (2.2). Panel B shows the separate parts of Equation (2.5) as well as the composition effect described in Equation (2.4). Panel C presents the results from decomposing the change in the within occupation gap following Equation (2.6). Values $\times 100$. Full-time working 25-54 year old West German men and women.

The long term effects mask variation between different episodes, however, with the largest effects taking place between 1985 and 1993. In fact, changes in skill prices were only beneficial for women's relative wages in that period but became small and even harmful for a decline in the wage gap afterwards. The opposite is true for skill changes which raised the wage gap between 1985 and 1993. This is also the time in which female labor force participation increased most. In fact, the share of women in the labor force rose by 3 percentage points during that episode. Hence, an explanation for the declining relative female skills could be the comparably large amount of new female labor market entrants with little experience.

Interestingly, similar timing results were found by Yamaguchi (2018, Table 12, p. 62) for the US: the return to motor skills declined up to 2000 inducing male and female wages to converge on average. After 2000, motor skill returns stayed roughly constant with no further effect on the gender wage gap. In contrast, changes in women's cognitive skills only had a small impact on the gap during 1980 – 1990 but induced a convergence of male and female wages after 2000.

In the last step, I decompose changes in the within wage gap (i.e., the unexplained part) further. I distinguish occupations in which male occupation shares $P_t(k|m)$ increased more than female occupation shares $P_t(k|f)$. This separates occupations which became more important for male employment over

time from occupations which became more important for female employment:

$$\underbrace{\Delta[\bar{e}_t^m - \bar{e}_t^f]}_{\text{within gap}} = \underbrace{\sum_{k=1}^K \Delta[\bar{e}_t^m - \bar{e}_t^f] \mathbb{1}(\Delta p_{k,t}^m > \Delta p_{k,t}^f)}_{\text{becoming more important for men}} + \underbrace{\sum_{k=1}^K \Delta[\bar{e}_t^m - \bar{e}_t^f] \mathbb{1}(\Delta p_{k,t}^m \leq \Delta p_{k,t}^f)}_{\text{becoming more important for women}} \quad (2.6)$$

Interestingly, the small overall residual effect of -0.06 log points hides that there have been converging wages between the sexes within occupations that became more important for male relative to female employment ($\Delta p_{k,t}^m > \Delta p_{k,t}^f$). This effect is substantial contributing -2.80 log points to the declining gap and, hence, comparable in magnitude to the overall wage structure effect. In contrast, however, male and female wages diverged by 2.74 log points within occupations which became more important for female employment over time. The bottom panel in Appendix Table 2.5 shows that male and female wages especially diverged in professional occupations contributing 1.44 log points to a growing gender wage gap whereas wages converged most within producing occupations leading to a decline of the wage gap by 2.29 log points. A possible explanation for this could be skill selection effects similar to the findings in Böhm et al. (2019). If many more women than men enter an occupation over time, they are likely to be less skilled than the incumbents partly because entrants are typically younger and therefore less skilled. The opposite effect might be true for occupations in which female employment declined more over time than male employment, with only the most skilled women remaining.

2.3.2 Proportion of Women Along the Wage Distribution

The influence of changing skill prices and changing occupational employment might be very different along the wage distribution. Whereas changing prices may be important for the decline in the average wage gap, the gap at the top or bottom of the wage distribution might be completely unaffected; for instance, because the distribution of occupational employment is very different at low and high wage percentiles; with workers in professional and managerial occupations at the top and workers in producing occupations located at middle and low percentiles (see Appendix Figure 2.6). For this reason, Table 2.4 decomposes the proportion of women at different percentiles of the combined male and female wage distribution into skill price and composition effects.

In total, the share of women in the full-time working labor force aged 25–54 increased from 30.4% in 1985 to 33.42% in 2010. Essentially, all of this rise took place during 1985–1993. The share of women differs strongly across different parts of the wage distribution, however. Whereas only 11.85% of workers at the 85th percentile were women in 1985, the proportion was 60.87%

2.3. IMPACT OF CHANGING SKILL RETURNS ON THE GENDER WAGE GAP 171

Table 2.4: Decomposition of proportion of women by percentile

	Level	Difference			
	1985	1993 - 1985	2001 - 1993	2010 - 2001	2010 - 1985
Proportion of women in sample	30.40	3.37	-0.48	0.13	3.02
<i>Panel A</i>					
Proportion of women at 85th percentile	11.85	5.60	1.88	0.83	8.30
<i>Scenarios relative to observed proportion</i>					
$\Delta \hat{\pi}_{k,t} = 0$		-2.57	0.48	0.48	-1.61
$\Delta \hat{s}_{k,t} = 0$		0.67	0.22	0.68	1.57
Rewgt. occupations		-0.71	0.61	-0.87	-0.97
Rewgt. occupations + $\Delta \hat{\pi}_{k,t} = 0$		-2.62	-2.01	-0.01	-4.64
<i>Panel B</i>					
Proportion of women at 50th percentile	20.10	4.96	1.90	3.83	10.68
<i>Scenarios relative to observed proportion</i>					
$\Delta \hat{\pi}_{k,t} = 0$		0.76	-0.40	-1.16	-0.80
$\Delta \hat{s}_{k,t} = 0$		-0.23	1.95	-3.29	-1.58
Rewgt. occupations		-1.87	-0.81	0.62	-2.06
Rewgt. occupations + $\Delta \hat{\pi}_{k,t} = 0$		0.33	-1.29	-1.10	-2.07
<i>Panel C</i>					
Proportion of women at 15th percentile	60.87	-3.95	-4.49	-2.65	-11.10
<i>Scenarios relative to observed proportion</i>					
$\Delta \hat{\pi}_{k,t} = 0$		5.31	-2.92	-0.70	1.70
$\Delta \hat{s}_{k,t} = 0$		3.35	-3.53	-1.94	-2.12
Rewgt. occupations		3.96	1.44	-3.12	2.28
Rewgt. occupations + $\Delta \hat{\pi}_{k,t} = 0$		8.06	-0.37	-1.13	6.55

Notes: The panels show the proportion of women in 0.5% bins around the 85th, 50th, and 15th percentile of the combined male and female wage distribution. The scenarios show counterfactual changes in proportions of women by percentile relative to the observed change. Hence, adding the observed change to the relative counterfactual numbers gives the counterfactual change in the proportion of women by percentile. The scenarios are as follows: $\hat{\pi}_{k,t} = 0$: subtract price growth between t and 1985 from wages observed in t ; $\hat{s}_{k,t} = 0$: subtract average occupational skill growth between t and 1985 from wages observed in t ; *Rewgt. occupations*: reweight occupations in t to match the occupation structure of 1985; *Rewgt. occupations + $\hat{\pi}_{k,t} = 0$* : reweight occupations in t to match the occupation structure of 1985, additionally subtract price growth between t and 1985 from wages observed in t . Values $\times 100$. Full-time working 25-54 year old West German men and women.

at the 15th percentile. Hence, women were both underrepresented at the top as well as overrepresented at the bottom of the wage distribution; compared to their overall representation in the work force. Over time, male and female wages converged because of both a rising share of women at the 85th percentile amounting to 20.15% in 2010 as well as a declining share at the 15th percentile resulting in a proportion of 49.77% in 2010. The share of women at the median was 20.10% in 1985 and increased over time up to 30.78%; so that by 2010, the representation of women at the median was almost equal to the representation of women in the overall sample.

To investigate the importance of changing prices on the change in female proportions by percentile, I perform the following exercise: take wages in year t and subtract the estimated change in skill prices between 1985 and t . Then, recalculate the proportion of women in this counterfactual wage distribution.

The results show that without skill prices having changed and everything else equal, the share of women at the 85th percentile in 2010 would have been

1.61 percentage points lower than observed as well as 1.7 percentage points higher at the 15th percentile.¹² This shows that changing prices led to less gender inequality across the distribution although the effects are modest. The women benefiting price changes were partly offset by changing skills as shown before for changes in the average gender wage gap. The scenario $\bar{\hat{s}}_{k,t} = 0$ shows the results when deducting average skill growth in occupation k between 1985 and t from wages observed in t . In fact, if skills had not changed, the proportion of women at the 85th percentile would have been 1.57 percentage points higher than observed. Roughly 40% of this dampening effect took place during 1985–1993 consistent with the large increase of inexperienced women into the labor force during that time. In line, the share of women at the 15th percentile would have declined by 2.12 percentage points more than observed if skills had not changed. Changing skills therefore contributed to a rise in gender wage inequality in all parts of the wage distribution. Hence, this effect has partly offset the women benefiting skill price changes.

Next, I investigate the influence of changing occupational employment for the proportion of women along the wage distribution. For that, I reweight observations following DiNardo et al. (1996) to estimate how the shares of women by percentile would have looked like if the occupational distribution of both men and women had not changed over time. The proportion of women would have been lower at the 85th percentile (0.97 percentage points) and higher at the 15th percentile (2.28 percentage points) if occupational employment patterns had not changed. This is mainly because women increasingly started to move into high wage occupations over time.

Last, I combine the experiments on skill price changes and employment changes by reweighting the skill price deducted wages to match the 1985 employment composition of men and women. This exercise shows that the proportion of women at the 85th percentile would have been 4.64 percentage points lower in 2010 whereas the share at the 15th percentile would have been 6.55 percentage points higher. Compared to an overall increase of the female share of 8.30 percentage points at the 85th percentile and a decrease of 11.10 points at the 15th percentile, the influence of the combination of changing employment and changing skill prices seems substantial, therefore. In summary, if skill prices and employment had not changed over the years, the proportion of women at the 85th percentile would have risen by 50% less compared to the observed increase. In line with that, the share of women at the 15th percentile would have declined by 55% less. Hence, changing skill prices and changing occupational employment patterns led to a decrease in gender wage inequality in both the upper and lower part of the wage distribution.

¹²Of course, a different evolution of skill prices would also have induced a different selection of men and women into the labor force and into occupations. In this paper, I abstract from that possibility. See Hsieh et al. (2019) for a structural approach which incorporates general equilibrium responses.

2.4 Conclusion

Starting in the 1980s, gender wage inequality has decreased in almost all developed countries (Olivetti and Petrongolo, 2016; Blau and Kahn, 2017). This paper has revisited the old question as to what extent changes in returns to occupational skill as well as changes in the occupational employment distribution were responsible for declining inequality between men and women.

In a first step, I found that almost the complete decrease of the gender wage gap can be attributed to changes in both occupational wages and the movement of women into high wage, managerial, and professional occupations. A simple decomposition into wage and composition effects showed that roughly 35% of the declining wage gap are explained by changing average occupational wages; with the remaining 65% being explained by changing employment patterns between the genders.

In the next step, I then made use of estimates for selection corrected skill price changes which hold constant the shifting skill composition by occupation over time. This accounts for the fact that the skill of the average worker in an occupation is endogenous to changing skill prices (Böhm et al., 2019); and hence, changes in average wages (wage structure effect) do not reflect the causal influence of changing prices on movements in the gender wage gap. Decomposing the wage structure effect further by means of these quality adjusted wages showed that women's wages profited both from declining prices in manufacturing occupations (employing large shares of the male workforce) as well as rising prices in sales and office occupations which are important for the female employment distribution.

Overall, changes in average skill by gender contributed to an increase of the wage gap during the last half of the 1980s and early 1990s. The main reason for that is the large rise in female labor force participation during that time, coming in hand with many inexperienced women entering the labor market. Changes in skill prices, in contrast, were beneficial for women's relative wages in the late 1980s. This trend reversed after mid 1990 in line with the findings of Yamaguchi (2018). In summary, I found that the effect attributed to changes in average wages is smaller than the causal effect that can be explained by changing skill prices. Therefore, the response of workers' skills to changing skill returns might have been the reason for the small (or even negative) wage structure effects in much of the previous literature concerned with explanations for why the gender wage gap declined (e.g., Blau and Kahn, 1997; Bacolod and Blum, 2010).

After having shown the importance of changes in skill returns as well as male and female occupational choices, an interesting topic for future research would be the question about why women increasingly started to move into professional and managerial occupations over time. Reasons for this could be occupation dependent changes in demand for female work and expanding possibilities to work part-time (Heathcote et al., 2010), lower frictions with respect to occupational choice including discrimination (Hsieh et al., 2019), or supply side

changes that made it possible for women to work longer hours; for instance, because of changes in fertility over time (Bailey et al., 2012; Cortés and Pan, 2019; Wasserman, 2019). In addition, investigating why male and female wages converged in only some (mostly low paying) but not all occupations would be important to understand why the level of gender wage inequality is still large and persistent with convergence having ceased since the early 2000s (Goldin, 2014; Goldin and Katz, 2016). The finding that gender inequality increased in occupations which became more important for female employment than male employment is suggestive that selection effects might also be important in that respect.

A Additional Tables

Table 2.5: Gender wage gap decomposition by professions

	Level	Difference			
	1985	1993 - 1985	2001 - 1993	2010 - 2001	2010 - 1985
Observed gap	44.24	-4.66	-2.41	-1.34	-8.42
Pred. gap, occupation \times year	20.45	-3.73	-3.13	-1.50	-8.37
Price structure		-4.96	0.57	0.35	-4.04
Mgr		0.10	-0.05	-0.02	0.03
Prof		1.25	-0.06	-0.21	0.97
Tech		-0.16	0.02	0.01	-0.13
Sales		-0.69	0.05	0.03	-0.62
Office		-3.76	0.27	0.20	-3.29
Prod		-0.67	0.00	0.02	-0.65
Op		-0.54	0.02	0.03	-0.49
Crafts		-1.73	0.20	0.13	-1.41
Svc		0.07	0.04	-0.02	0.09
Care		1.19	0.08	0.18	1.45
Skill structure		3.52	-1.54	-0.89	1.08
Mgr		0.10	0.06	0.08	0.24
Prof		-1.04	0.31	0.52	-0.21
Tech		0.73	0.15	-0.07	0.81
Sales		-0.20	-0.27	0.10	-0.37
Office		0.79	-1.18	-0.52	-0.91
Prod		2.07	0.10	-0.07	2.09
Op		1.35	-0.53	-0.98	-0.17
Crafts		3.21	-0.50	-0.77	1.94
Svc		-0.82	0.26	0.45	-0.11
Care		-2.66	0.06	0.38	-2.23
Composition effect		-2.29	-2.17	-0.95	-5.41
Mgr		-5.01	-6.26	-2.31	-13.58
Prof		-5.35	-3.68	-12.61	-21.64
Tech		-7.43	-3.34	-2.93	-13.70
Sales		-7.10	1.98	6.73	1.61
Office		4.65	15.64	16.02	36.31
Prod		12.89	-1.20	4.33	16.02
Op		-2.01	1.38	0.82	0.19
Crafts		7.37	-12.56	-9.35	-14.54
Svc		16.09	12.86	11.60	40.55
Care		-16.39	-7.00	-13.25	-36.64
Within occupation gap	23.79	-0.93	0.72	0.16	-0.06
Mgr		0.03	0.20	0.46	0.69
Prof		0.48	0.48	0.48	1.44
Tech		0.11	-0.06	-0.09	-0.05
Sales		0.05	0.03	-0.12	-0.04
Office		-0.64	0.62	0.76	0.75
Prod		-0.74	-0.55	-0.99	-2.29
Op		0.13	0.11	-0.01	0.22
Crafts		-0.14	-0.18	-0.32	-0.64
Svc		-0.27	-0.11	-0.03	-0.41
Care		0.06	0.18	0.03	0.27

Notes: The panels show the results from Equation (2.2) and Equation (2.5) split by profession. Mgr: managers; Prof: professions; Tech: technicians; Prod: production; Op: Operators; Crafts: craftsmen; Svc: services. Classification of detailed occupations follows Acemoglu and Autor (2011), see Böhm et al. (2019) for the exact mapping. Values $\times 100$. Full-time working 25-54 year old West German men and women.

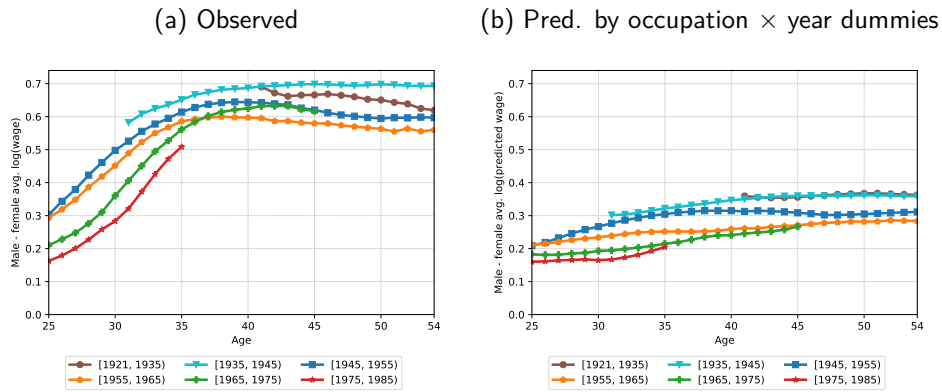
Table 2.6: Average gender wage gap decomposition including part-time workers

	Level	Difference			
	1985	1993 - 1985	2001 - 1993	2010 - 2001	2010 - 1985
<i>Panel A</i>					
Observed gap	59.09	-4.48	-3.57	0.27	-7.79
Pred. gap, occupation \times year	31.59	-3.82	-3.98	-0.38	-8.17
<i>Panel B</i>					
Wage structure effect		-0.79	-1.05	1.08	-0.76
Growing prices		-1.43	0.23	0.23	-0.98
Declining prices		-4.06	0.50	0.34	-3.22
Skill changes		4.70	-1.78	0.52	3.44
Composition effect		-3.03	-2.93	-1.46	-7.41
<i>Panel C</i>					
Within occupation gap	27.51	-0.67	0.40	0.65	0.38
$\Delta p_{k,t}^m > \Delta p_{k,t}^f$		-1.44	-0.83	-0.60	-2.86
$\Delta p_{k,t}^m \leq \Delta p_{k,t}^f$		0.77	1.23	1.24	3.24

Notes: Panel A shows log wages of men minus log wages of women for two different wage measures: observed wages and predicted wages from a regression of wages on 120 occupation dummies interacted with year dummies, i.e. the between occupation gap in Equation (2.2). Panel B shows the separate parts of Equation (2.5) as well as the composition effect described in Equation (2.4). Panel C presents the results from decomposing the change in the within occupation gap following Equation (2.6). Values $\times 100$. Full-time and part-time working 25-54 year old German men and women. The wage of a part-time worker is not adjusted for hours. Instead, the wage refers to a working day as is also the case for full-time workers.

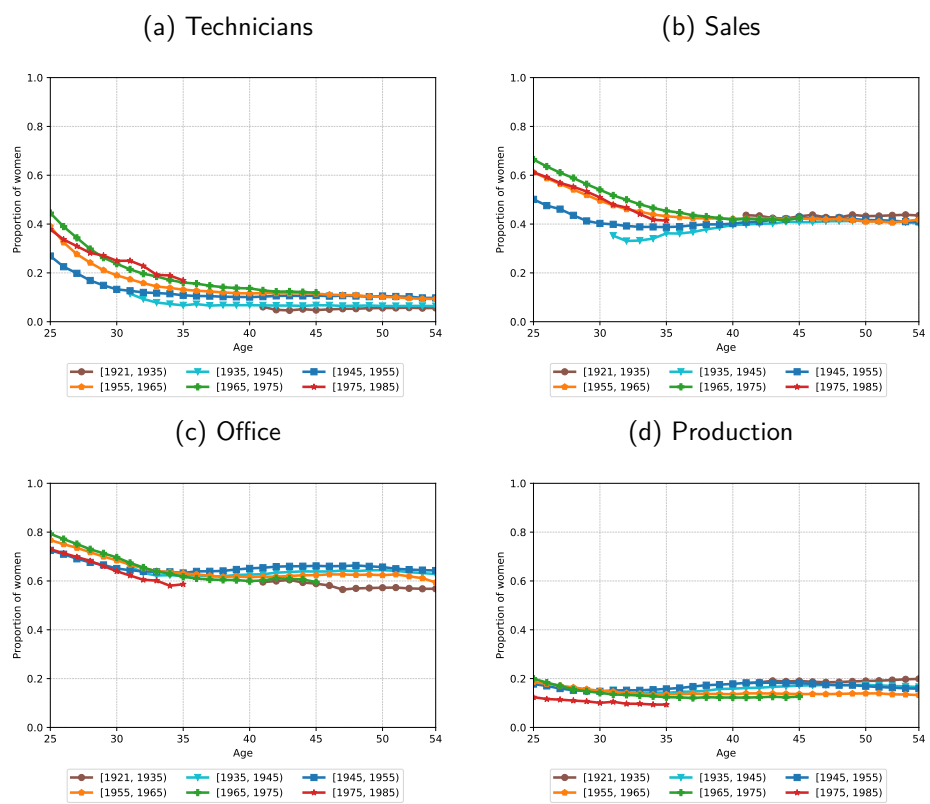
B Additional Figures

Figure 2.4: Wage gap by cohort including part-time workers



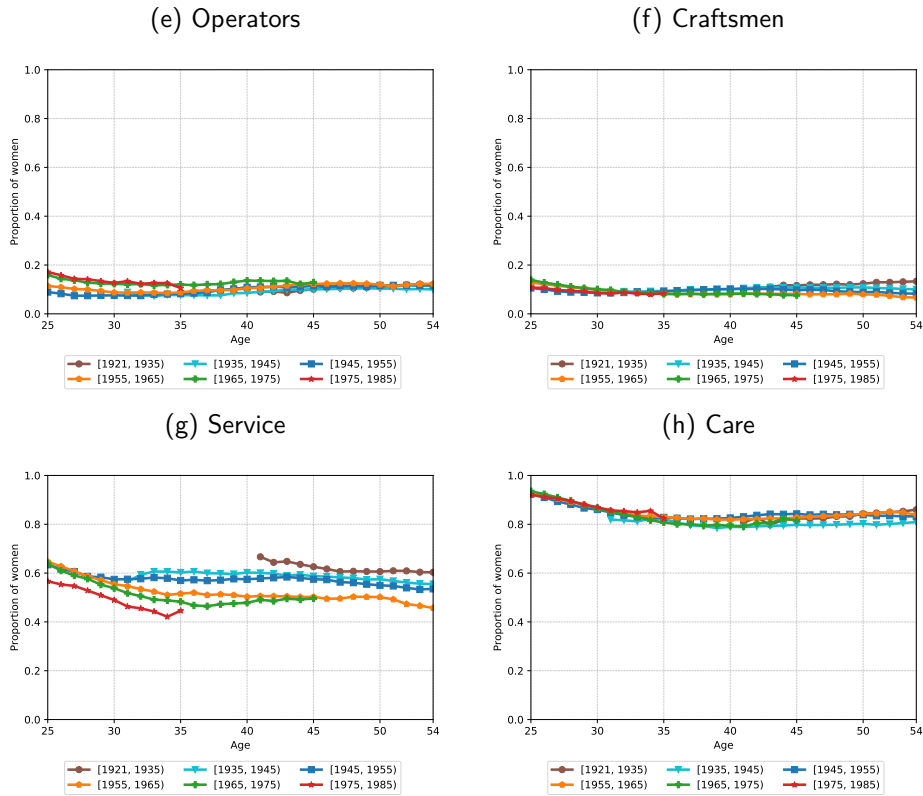
Notes: Figure 2.4a shows log wages of men minus log wages of women by cohort and age. Figure 2.4b repeats this but uses the predictions from a regression of log wages on occupation times year dummies to calculate the wage gap. Full-time and part-time working 25-54 year old German men and women. The wage of a part-time worker is not adjusted for hours. Instead, the wage refers to a working day as is also the case for full-time workers.

Figure 2.5: Proportion of women in professions by cohort



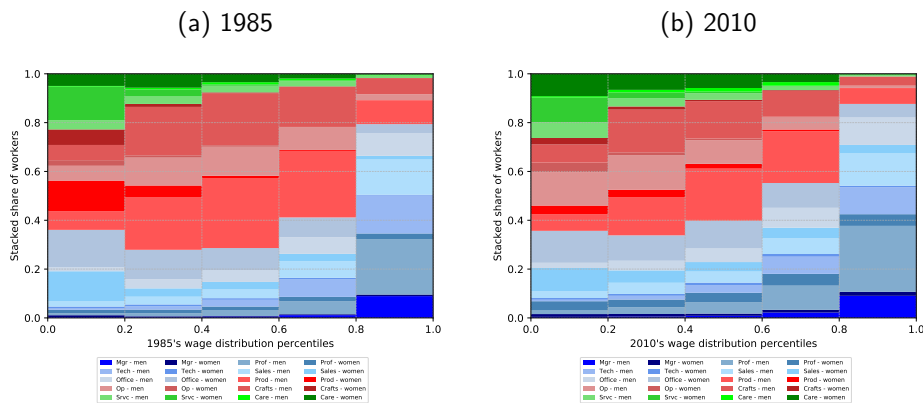
Notes: Classification of detailed occupations follows Acemoglu and Autor (2011), see Böhm et al. (2019) for the exact mapping. Full-time working 25-54 year old West German men and women.

Figure 2.5: Proportion of women in professions by cohort



Notes: Classification of detailed occupations follows Acemoglu and Autor (2011), see Böhm et al. (2019) for the exact mapping. Full-time working 25-54 year old West German men and women.

Figure 2.6: Profession \times gender combinations along the wage distribution



Notes: Each rectangle is proportional to the share of workers in the respective occupation times gender bin. Wage distribution refers to the combined male and female wage distribution. Mgr: managers; Prof: professions; Tech: technicians; Prod: production; Op: Operators; Crafts: craftsmen; Svc: services. Classification of detailed occupations follows Acemoglu and Autor (2011), see Böhm et al. (2019) for the exact mapping. Full-time working 25-54 year old West German men and women.

References

- Acemoglu, Daron and David H. Autor (2011). "Skills, Tasks and Technologies: Implications for Employment and Earnings". In: *Handbook of Labor Economics, vol. 4 Part B*. Edited by David Card and Orley Ashenfelter. Elsevier. Chapter 12, pages 1043–1171.
- Adda, Jérôme, Christian Dustmann, and Katrien Stevens (2017). "The Career Costs of Children". In: *Journal of Political Economy* 125.2, pages 293–337.
- Altonji, Joseph G. and Rebecca M. Blank (1999). "Race and Gender in the Labor Market". In: *Handbook of Labor Economics* 3, pages 3143–3259.
- Bacolod, Marigee P. and Bernardo S. Blum (2010). "Two Sides of the Same Coin US "Residual" Inequality and the Gender Gap". In: *Journal of Human Resources* 45.1, pages 197–242.
- Bailey, Martha J, Brad Hershbein, and Amalia R Miller (2012). "The Opt-in Revolution? Contraception and the Gender Gap in Wages". In: *American Economic Journal: Applied Economics* 4.3, pages 225–54.
- Beaudry, Paul and Ethan Lewis (2014). "Do Male-Female Wage Differentials Reflect Differences in the Return to Skill? Cross-City Evidence from 1980–2000". In: *American Economic Journal: Applied Economics* 6.2, pages 178–94.
- Black, Sandra E. and Alexandra Spitz-Oener (2010). "Explaining Women's Success: Technological Change and the Skill Content of Women's Work". In: *Review of Economics and Statistics* 92.1, pages 187–194.
- Blau, Francine D. and Lawrence M. Kahn (1997). "Swimming Upstream: Trends in the Gender Wage Differential in the 1980s". In: *Journal of Labor Economics* 15.1, pages 1–42.
- Blau, Francine D. and Lawrence M. Kahn (2017). "The Gender Wage Gap: Extent, Trends, and Explanations". In: *Journal of Economic Literature* 55.3, pages 789–865.
- Blau, Francine D, Peter Brummund, and Albert Yung-Hsu Liu (2013). "Trends in Occupational Segregation by Gender 1970–2009: Adjusting for the Impact of Changes in the Occupational Coding System". In: *Demography* 50.2, pages 471–492.
- Böhm, Michael, Hans-Martin von Gaudecker, and Felix Schran (2019). "Occupation Growth, Skill Prices, and Wage Inequality". Working Paper.
- Borghans, Lex, Bas Ter Weel, and Bruce A. Weinberg (2014). "People Skills and the Labor-Market Outcomes of Underrepresented Groups". In: *ILR Review* 67.2, pages 287–334.
- Card, David and John E. DiNardo (2002). "Skill-Biased Technological Change and Rising Wage Inequality: Some Problems and Puzzles". In: *Journal of Labor Economics* 20.4, pages 733–783.
- Cortes, Guido Matias (2016). "Where Have the Middle-Wage Workers Gone? A Study of Polarization Using Panel Data". In: *Journal of Labor Economics* 34.1, pages 63–105.

- Cortes, Patricia and Jessica Pan (2018). "Occupation and gender". In: *The Oxford Handbook of Women and the Economy*. Edited by Marvin Kusters. Oxford University Press, pages 425–452.
- Cortés, Patricia and Jessica Pan (2019). "When Time Binds: Substitutes for Household Production, Returns to Working Long Hours, and the Skilled Gender Wage Gap". In: *Journal of Labor Economics* 37.2, pages 351–398.
- DiNardo, John, Nicole M. Fortin, and Thomas Lemieux (1996). "Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semiparametric Approach". In: *Econometrica* 64.5, pages 1001–1044.
- Duncan, Otis Dudley and Beverly Duncan (1955). "A Methodological Analysis of Segregation Indexes". In: *American Sociological Review* 20.2, pages 210–217.
- Fuchs, Victor R. (1975). "A Note on Sex Segregation in Professional Occupations". In: *Explorations in Economic Research*. Volume 2. NBER, pages 105–111.
- Goldin, Claudia D. (1990). *Understanding the Gender Gap: An Economic History of American Women*. Oxford University Press.
- Goldin, Claudia D. (2014). "A Grand Gender Convergence: Its Last Chapter". In: *American Economic Review* 104.4, pages 1091–1119.
- Goldin, Claudia D. and Lawrence F. Katz (2016). "A Most Egalitarian Profession: Pharmacy and the Evolution of a Family-Friendly Occupation". In: *Journal of Labor Economics* 34.3, pages 705–746.
- Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante (2010). "The Macroeconomic Implications of Rising Wage Inequality in the United States". In: *Journal of Political Economy* 118.4, pages 681–722.
- Hsieh, Chang-Tai, Erik Hurst, Charles I. Jones, and Peter J. Klenow (2019). "The Allocation of Talent and U.S. Economic Growth". Working Paper, forthcoming in: *Econometrica*.
- Juhn, Chinhui, Kevin M. Murphy, and Brooks Pierce (1991). "Accounting for the Slowdown in Black-White Wage Convergence". In: *Workers and Their Wages*. Edited by Marvin Kusters. AEI Press, pages 107–143.
- Juhn, Chinhui, Kevin M. Murphy, and Brooks Pierce (1993). "Wage Inequality and the Rise in Returns to Skill". In: *Journal of Political Economy* 101.3, pages 410–442.
- Kleven, Henrik, Camille Landais, and Jakob E. Søgaa (2018). "Children and Gender Inequality: Evidence from Denmark". Working Paper, forthcoming in: *American Economic Journal: Applied Economics*.
- Mulligan, Casey B. and Yona Rubinstein (2008). "Selection, Investment, and Women's Relative Wages Over Time". In: *Quarterly Journal of Economics* 123.3, pages 1061–1110.
- Neumark, David (1988). "Employers' Discriminatory Behavior and the Estimation of Wage Discrimination". In: *Journal of Human Resources* 23.3, pages 279–295.

- Ngai, Rachel L. and Barbara Petrongolo (2017). "Gender Gaps and the Rise of the Service Economy". In: *American Economic Journal: Macroeconomics* 9.4, pages 1–44.
- OECD (2019). *Employment: Share of Employed in Part-Time Employment, by Sex and Age Group*. data retrieved from OECD.Stat on 19.07.2019. URL: <https://stats.oecd.org/index.aspx?queryid=54746#>.
- Olivetti, Claudia (2006). "Changes in Women's Hours of Market Work: The Role of Returns to Experience". In: *Review of Economic Dynamics* 9.4, pages 557–587.
- Olivetti, Claudia and Barbara Petrongolo (2016). "The Evolution of Gender Gaps in Industrialized Countries". In: *Annual Review of Economics* 8.1, pages 405–434.
- Roy, A. D. (1951). "Some Thoughts on the Distribution of Earnings". In: *Oxford Economic Papers* 3.2, pages 135–146.
- Smith, James and Finis R. Welch (1989). "Black Economic Progress after Myrdal". In: *Journal of Economic Literature* 27.2, pages 519–64.
- Wasserman, Melanie (2019). "Hours Constraints, Occupational Choice, and Gender: Evidence from Medical Residents". Working Paper.
- Weinberg, Bruce A (2000). "Computer Use and the Demand for Female Workers". In: *ILR Review* 53.2, pages 290–308.
- Weisskoff, Francine D. Blau (1972). "'Women's Place' in the Labor Market". In: *American Economic Review* 62.1, pages 161–166.
- Welch, Finis (2000). "Growth in Women's Relative Wages and in Inequality among Men: One Phenomenon or Two?" In: *American Economic Review: Papers and Proceedings* 90.2, pages 444–449.
- Yamaguchi, Shintaro (2018). "Changes in Returns to Task-Specific Skills and Gender Wage Gap". In: *Journal of Human Resources* 53.1, pages 32–70.

Chapter 3

Locational Choice and Spatial Wage Inequality

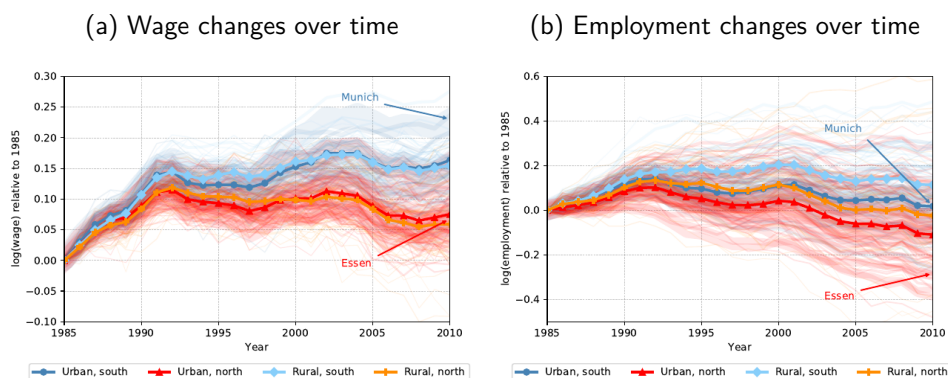
3.1 Introduction

Inequality in developed countries has increased substantially over the last decades (Acemoglu and Autor, 2011). One important factor contributing to this trend are increasing spatial disparities between booming and declining areas like San Francisco and Detroit in the US or Munich and Essen in Germany (Blanchard and Katz, 1992; Moretti, 2012; Dauth, Findeisen, Moretti, et al., 2019). In fact, these growing disparities have led many countries to implement “place based policies” aimed at reducing regional wage inequality (Glaeser and Gottlieb, 2008; Kline and Moretti, 2014). The reasons for why wages and employment in some areas are rising while shrinking in others are not yet fully understood, however.

Germany is no exception with respect to large heterogeneity in wage and employment growth across space. While real wages of prime age men in Germany rose by around 10 log points between 1985 and 2010 on average, this wage growth was very unevenly distributed geographically.¹ Figure 3.1a shows average log wages of 182 distinct German local labor markets relative to 1985. Wages in Southern German local labor markets increased by up to 28 log points but much less so in northern regions – irrespective of being rural or urban. On average, southern wages grew by roughly 16 log points and by only 7 log points in the north. The same pattern is to be found in Figure 3.1b with respect to changes in regional employment. Employment increased in Southern Germany and decreased in the north.

¹One may view it also as a consequence of these diverging trends that the German government has appointed a commission on “equitable living conditions” in 2018. The aim of this commission is to provide recommendations for achieving spatial convergence in factors ranging from local debt and employment opportunities to infrastructure, public services and social cohesion. The federal state of Bavaria has even made it the state’s obligation to ensure equitable living standards across rural and urban places by incorporating that obligation into its constitution.

Figure 3.1: Changes in employment and wages across West German regions



Notes: Figure 3.1a shows the change in average log wages within local labor markets over time. Figure 3.1b shows the change in log employment. Shaded lines in the background represent 182 West German detailed local labor markets. The four groups show a classification of local labor markets into northern and southern, rural and urban areas (see Appendix Figure 3.10). Southern areas are located in Baden-Wuerttemberg and Bavaria. The classification of local labor markets into rural and urban is based on BBSR (2014). The thickness of a shaded background line corresponds to the number of employed workers in a local labor market averaged across years 1985 until 2010.

There exist many potential explanations for this but the simultaneous rise in southern wages and employment suggests that demand for workers' skills increased relatively more in the south than in the north. Another potential explanation would be that worker quality rose in Southern Germany; or more generally, that worker quality rose within regions where employment increased. This paper attempts to separate these two competing explanations with the aim to get a more comprehensive answer to the question: are increasing geographical differences driven by spatial changes in the demand for skill or the supply of skill?

For that, I estimate how regional wage premia paid for a unit of skill have changed over time in detailed local labor markets by exploiting administrative panel data with high quality information on workers' location choices. Importantly, the employed estimation method allows workers with heterogeneous skills to self select into local labor markets depending on where they can make the highest wages and profit from other, non-wage aspects which might influence their utility in a generalized Roy model framework (Roy, 1951; Dahl, 2002). This accounts for the fact that differential changes in wages across local labor markets can either be the result of changes in *wage premia* paid to all workers within a local labor market, or be the result of changes in the *quality* of local workers.

I adopt a method which originally focused on estimating returns to occupational skills in the face of worker self selection (Böhm et al., 2019). Therefore, I exploit accelerations and decelerations of individual workers' wage growth over time and by region together with workers' location decisions to estimate changes

in local wage premia.² Intuitively, the estimated wage premium within a region is rising if workers' individual wage growth accelerated over time and is declining vice versa. The estimating equation is derived from an explicit model of locational choice based on workers' comparative advantage which allows for a structural interpretation of the estimates. This is in contrast to more common fixed effect models (Abowd et al., 1999; Combes, Duranton, and Gobillon, 2008; Combes, Duranton, Gobillon, and Roux, 2012; Dauth, Findeisen, Moretti, et al., 2019). Importantly, identification in the employed method does not rely on the exogenous mobility assumption (made explicit in Combes, Duranton, and Gobillon, 2008) and so is largely robust to workers switching location because of idiosyncratic error realizations. Apart from that, one can easily control for differences in life cycle skill accumulation profiles across regions (see also the discussion about the importance of this in Glaeser, 1999; Wheeler, 2006; De la Roca and Puga, 2017).

By exploiting the conditional wage growth of regional movers, I am also able to control and estimate how other, non-wage aspects like local amenities have differentially changed over time. The idea for this is that (conditional) average wage growth associated with switching from a southern region to a northern region should be relatively high if amenities in the north are low. Wage growth then serves as a compensating differential (Rosen, 1986; Roback, 1988; Sorkin, 2018).

I find that local labor markets with growing employment did not feature rising wage premia. This conflicts with the fact that wages in growing regions increased more than in regions where employment declined. It suggests that the simultaneous rise in southern wages and employment is rather due to sorting of high skilled workers into the south than due to increasing southern demand. In fact, the residual between growth in average wages and changes in (quality adjusted) wage premia is informative about changes in worker quality within the model framework. I find that the quality of workers rose more in regions with growing employment, i.e. Southern Germany. Put differently, this suggests that regions with expanding employment exhibit positive skill selection.³

There exist several potential explanations for the positive correlation between worker quality changes and employment growth which are independent of the model. First, shifts in observables such as the occupation, education and industry structure might have been different between rising and deteriorating regions (Findeisen and Südekum, 2008). Second, increases in employment density might have raised agglomeration economies making workers and firms become

²See Wheaton (1979) and Blanchard and Katz (1992) for early papers which attempt to identify the role of regional supply and demand shocks from local employment growth and wages.

³This is in strong contrast to the findings in Böhm et al. (2019) regarding occupational growth. Böhm et al. (2019) find that occupational wage and employment growth are uncorrelated because of selection effects in Germany. This is similar to findings for the US by Hsieh et al. (2019).

more productive in booming regions (Duranton and Puga, 2004). Third, skilled workers could increasingly sort into high amenity places (Diamond, 2016). I cannot completely separate these (non-exhaustive) explanations. Instead, I will focus on the importance of changes in the occupation, education and industry structure as well as local amenities in greater depth.

I find that regions, in which worker quality improved, also experienced disproportionate increases in the amount of workers employed in managerial, professional and technical (high wage, high skill) occupations. Further, growth in high-skill service and machine building industries as well as growth in the number of university graduates was larger in these regions. In contrast, worker quality declined in markets with rising production and manufacturing employment.

To quantify the magnitude of changes in the industry, education and occupation structure for the correlation between worker quality and employment changes, I perform a reweighting exercise following DiNardo et al. (1996). I reweight observations in 1985 by assigning higher weights to occupation-region (as well as industry-occupation-region and education-region) cells if the number of workers within such a cell increased over time. Using these weights, I recalculate the elasticities between worker quality changes and employment changes. I find that the elasticities drop substantially when reweighting with occupation and industry but much less so for education. This underscores the importance of spatially uneven occupational and industrial change for geographic differences in wage growth (Duranton, 2007; Findeisen and Südekum, 2008; Südekum, 2008; Amior and Manning, 2018).

Furthermore, I find that regions with high estimated amenity values have experienced positive worker quality growth. This is consistent with skilled workers increasingly sorting into areas with high living standards (Diamond, 2016). As the amenity estimates are somewhat noisy because they are identified from movers' wage growth, I additionally collected data on an amenity index reflecting amenities such as sunshine duration, restaurant and physician density as well as data on crime rates and the number of students per school. The amenity index alone is highly predictive of worker quality changes. In line with that, worker quality also increased in areas with low crime rates as a proxy for disamenities. When taking all model independent explanatory variables together (i.e. occupation, industry, education and amenity proxies), 59% of the spatial variation in quality changes can be explained by these observables which is sizable.

In the last step, I use the estimates of changes in wage premia and worker quality to shed light on the rise in spatial wage inequality: I investigate the importance of changing worker quality and changing wage premia for the rising density wage premium; a measure which reflects inequality between rural and urban places and so is important for overall spatial inequality.

I find that dense areas experienced a larger increase in worker quality consistent with the rising density premium. However, improving worker quality was completely offset by declining wage premia in heavily populated areas. Instead, the reason for the rise in the density premium is due to a decline in employment

of the formerly densest places (industrial places like the Ruhr area) and a contemporaneous increase in the density of formerly less populated places (mainly in rural Bavaria). Together with changes in average wages being unrelated to 1985's market densities, this has boosted the density wage premium overall.

The remainder of this chapter is as follows. Section 3.2 discusses the related literature. Section 3.3 shows facts of regional wage dispersion in Germany, and presents first evidence on workers' spatial self selection patterns. In Section 3.4, I lay out the employed method and give some evidence on the appropriateness of the required identifying assumptions. I describe the results on how wage premia and worker quality have changed across German local labor markets in Section 3.5. I also evaluate possible determinants of spatial differences in worker quality changes in that section. Section 3.6 analyzes reasons for the rising density wage premium. Section 3.7 concludes.

3.2 Connection to the Existing Literature

This paper is related to several strands in the urban economics literature. First and foremost, it is related to the literature about spatial sorting of skill across places and the urban wage premium (Berry and Glaeser, 2005; Moretti, 2013; Eeckhout et al., 2014; Dauth, Findeisen, Moretti, et al., 2019).

The closest paper in that part of the literature is work by Combes, Duranton, and Gobillon (2008) who investigate the importance of worker sorting on spatial wage disparities. Using French panel data from administrative records, they find that the elasticity of wages with respect to density drops substantially when including worker fixed effects. They interpret this finding as evidence of sorting of more skilled workers into denser areas.⁴

Three major points distinguish their work from this paper. First, Combes, Duranton, and Gobillon (2008) do not exploit the dynamics of sorting over time. Instead, they are interested in absolute wage differences between more or less dense places. Second, their approach requires much stricter assumptions on the determinants of location decisions. Importantly, they assume that moving decisions must not be based on realizations of idiosyncratic error draws. Although, the approach taken in this paper strictly needs this assumption as well, Böhm et al. (2019) show in Monte Carlo analyses that their method is much more robust to the violation of exogenous mobility. Third, Combes, Duranton, and Gobillon (2008) remain largely agnostic about the mechanism for the sorting effect as they do not derive their estimation strategy from an explicit model describing workers' decisions. Despite that, they acknowledge that self selection of workers based on their comparative advantage may be a reason for higher

⁴In follow up work, Combes, Duranton, Gobillon, and Rouxe (2012) then estimate full distributions of skill in local labor markets and find that dense places are not only more skilled on average but are also characterized by a higher dispersion of skill consistent with Eeckhout et al. (2014).

wages in denser places.

The second part of the literature this paper contributes to, deals with the importance of dynamic skill accumulation effects for differences in wages between urban and rural workers (Glaeser, 1999; D'Costa and Overman, 2014). For example, in contrast to Combes, Duranton, and Gobillon (2008), De la Roca and Puga (2017) find that workers in cities do not have higher unobserved ability than other workers initially, but instead accumulate skills at a faster rate through learning by working in dense places. Workers take the accumulated skills with them when leaving to less dense places. De la Roca and Puga (2017) argue that the sorting effect found by Combes, Duranton, and Gobillon (2008) is rather due to the higher skill accumulation present in cities which leads to biased fixed effects estimates. Complementary to their work, this paper also accounts for the importance of spatial differences in skill accumulation. I find that, indeed, denser places exhibit much steeper skill accumulation profiles.

Third, this paper adds to the literature on the importance of local amenities for location decisions, urban growth and inequality (Rosen, 1986; Roback, 1988; Glaeser, Kolko, et al., 2001). For instance, Diamond (2016) investigates the changing sorting pattern of college and non-college workers in the US based on wages, housing prices as well as local amenities. She finds that college workers increasingly concentrate because of differential changes in local demand. This sorting comes at the cost of higher increases in local rents in skilled areas, however. In turn, these high rent places also have high amenities to compensate and attract high skilled workers. In the estimation, I account for the fact that workers' location decisions not only depend on potential wages but also depend on local prices and amenities. In a final step, I exploit correlations between the estimated amenity values (as well as externally collected amenity data) and worker quality estimates.

The last part of the related literature deals with the role of spatial changes in the occupation and industry structure for local employment and wage growth (Blanchard and Katz, 1992; Duranton, 2007; Dustmann and Glitz, 2015; Amior and Manning, 2018). With respect to the German case, for instance, Findeisen and Südekum (2008) evaluate the importance of changes in the local industry structure as well as human capital for urban growth. In line with them, I also find that regions, which were able to "reinvent" themselves by adopting to a changing occupation-industry structure, exhibited rises in worker quality. This highlights the importance of structural change taking place unevenly across space (Autor, 2019).

3.3 Spatial Changes in Wages and Employment

3.3.1 Data

For all of the analyses, I make use of the Sample of Integrated Labor Market Biographies (SIAB) Regional File provided by the IAB Institute at the German

Federal Employment Agency. The SIAB is a 2% random sample of administrative social security records from 1975 to 2014. It is representative of 80% of the German workforce and includes employees covered by social security, marginal part-time employment, benefit receipts, officially registered as job-seeking or participating in programs of active labor market policy. The SIAB excludes the self-employed, civil servants, and individuals performing military service. Most notably, it contains individuals' full employment histories including detailed data on wages, place of work, occupation and industry along with socio-demographics such as age, gender, and the level of education.

Place of work corresponds to the political district (Landkreis) a plant is located in.⁵ I map every district (in former West Germany) into one of 182 local labor markets and drop all observations within East German districts. The classification of local labor markets is based on commuting flows of workers as described in Kropp and Schwengler (2011) and Dauth, Findeisen, Moretti, et al. (2019).⁶ I further aggregate these 182 local labor markets into four categories: urban - south, urban - north, rural - south, and rural - north. The assignment of regions to rural and urban markets is based on a classification of BBSR (2014). Southern regions are located in the federal states of Baden-Wuerttemberg and Bavaria. Appendix A.1 provides more information on the spatial distribution of this classification.

I prepare the data in the exact same way as Böhm et al. (2019). Importantly, I transfer the spell structure into a yearly panel by deleting all spells except for the longest spell within a year. I stochastically impute wages above the social security limit (as in Card et al., 2013), and restrict the sample to 25 to 54 year old German men working full-time in former West Germany between 1975 and 2010.

For some parts of the analysis, I rely on aggregations of detailed occupations, industries and education. See Appendix A for the details. Appendix A also contains information on collected amenity data and local price indices.

3.3.2 The Geography of Wage and Employment Changes Across Germany

German wage inequality has increased a lot during the last few decades (e.g., Dustmann, Ludsteck, et al., 2009). An important contributing factor are rising wage differentials across space (Moretti, 2013; Dauth, Findeisen, Moretti, et al., 2019). What are the drivers of these increasing spatial differences? Are

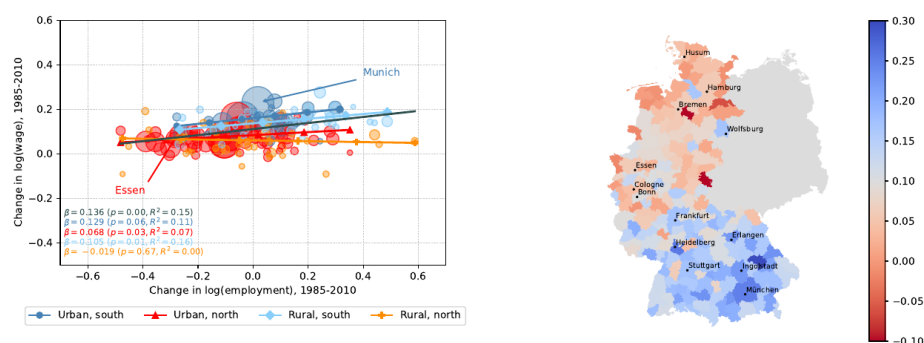
⁵Theoretically, a move between districts can therefore also reflect the move of a plant. This is very rare in practice, however.

⁶The number of 182 local labor markets differs from the 204 markets which Dauth, Findeisen, Moretti, et al. (2019) use in their main analysis because some rather sparsely populated districts are aggregated further due to data anonymization in the Scientific Use File I exploit. Hence, where the mapping is not one to one, I further aggregate these local labor markets into larger markets. See Ganzer et al. (2017, Table A.9) for the exact aggregation of districts because of anonymization in the SIAB Scientific Use File.

they the result of demand having declined in previous industrial clusters like the Ruhr area; but having increased in Southern Germany where many high tech firms operate (Findeisen and Südekum, 2008)? Or are these spatial changes the result of high skilled workers increasingly moving to high amenity areas (Diamond, 2016)?

Figure 3.2: Relation between employment growth and wage growth

(a) Employment changes and wage changes (b) Average wage changes, 1985 – 2010



Notes: The vertical axis in Figure 3.2a shows the change in average log wages between 1985 and 2010. The horizontal axis depicts the change in log employment between 1985 and 2010. One bubble represents one of 182 West German local labor markets aggregated from political districts as suggested by Kropp and Schwengler (2011). Bubble size corresponds to the number of employed workers in a market averaged across years 1985 to 2010. The four groups show a classification of local labor markets into northern and southern, rural and urban areas (see Appendix Figure 3.10). Southern areas are located in Baden-Wuerttemberg and Bavaria. The classification of local labor markets into rural and urban is based on BBSR (2014). Regression lines across all areas (black) and within the four broad groups (colored) are weighted by the number of employed workers. Figure 3.2b plots the change in average log wages between 1985 and 2010 across Germany. East Germany is left out of the analysis.

Figure 3.2a shows that regional wage and employment changes have been strongly correlated in West Germany. Rising employment was accompanied by rising wages: a region with 10% more employment growth experienced 1.4% more wage growth on average. Geographically, wages primarily grew in Southern Germany (see Figure 3.2b) and stayed largely constant or even declined in the north.

There are two distinct explanations for the positive correlation between wage and employment growth. First, this positive correlation seemingly suggests that regions with rising employment experienced increases in demand for skill leading to employment and wage gains. This assessment, however, crucially relies on the composition of worker quality not having changed over time in response to the change in demand. This brings me to the second possible explanation: if the quality of the average worker in a growing region increased much more than the quality of the average worker in a region with declining employment, changing supply of skill would be the more important factor for growing wage disparities.

To separate these explanations, the next section presents a method for es-

timating how wage premia, paid for a unit of skill, have changed over time and across space. Crucially, the method will account for the possibility that workers might self select into regions thereby altering the skill composition (Combes, Duranton, and Gobillon, 2008, p. 15, footnote 27); and that life cycle skill accumulation might be different across regions, for instance, because urban places facilitate learning (Glaeser, 1999; De la Roca and Puga, 2017).

But before that, I will briefly summarize evidence that the positive correlation between wage and employment growth is rather due to a changing skill composition than due to spatial differences in altering wage premia. Appendix B presents four stylized facts. It shows correlations between local employment growth and (1) changes in occupational employment shares, (2) changes in industrial employment shares, (3) changes in education shares as well as (4) relative wages of regional entrants.⁷ These correlations are informative about the observable types of attracted workers (e.g., high or low skilled) and the types of jobs workers in growing regions increasingly perform (e.g., service or high-tech). In total, these stylized facts suggest that growing regions were able to attract workers of comparably high quality as well as jobs performed in occupations and industries located in the upper part of the wage distribution.

In fact, the share of workers in managerial, professional and technical (high wage, high skill) occupations increased much more in growing than in declining regions (see Appendix B.1). The opposite is true for service and care (low wage, low skill) occupations whose share decreased in regions with growing employment. The geography of changes in the share of managerial, professional, and technical occupation is striking. Despite the number of workers in that group increased by roughly six percentage points between 1985 and 2010 overall, this rise almost exclusively took place in Southern Germany.

A similar result applies to changes in industrial employment shares (see Appendix B.2). I aggregate industries following Bárány and Siegel (2018) into manufacturing, high-skill services and low-skill services. In addition, I distinguish the machine building sector from the remaining manufacturing sectors because of Germany's large and increasing number of firms engaged in machine engineering (Findeisen and Südekum, 2008). In fact, wages within machine building are the highest across industries and have even increased much more over time than in high-skill services (21 log points compared to 8 log points). On average, the share of workers employed in machine engineering firms increased in growing local labor markets and declined vice versa. In contrast, both high- and low-skill services primarily grew in regions with declining employment. Therefore, not only the share of workers employed in high paying *occupations* rose within expanding regions over time; but also the share of workers employed in high paying *machine building* firms. Similar to the shares of high wage, managerial, professional and technical occupations, most of the local labor markets with increasing machine building shares were located in rather rural Southern

⁷Appendix B also shows maps which plot the variation of these variables across space.

Germany. This region underwent a dramatic transformation process. It was formerly “dominated by traditional industries but then developed rapidly over the 1970s and 1980s to become one of Germany’s leading high-tech states” (Findeisen and Südekum, 2008, p. 15).

Contrary to that, there is no relation between employment growth and changes in the share of university graduates (see Appendix B.3). Indeed, neither the share of university graduates nor the shares of medium or low educated workers increased differentially between growing and declining regions. The main reason is that growth in the proportion of university graduates primarily took place within cities which did not become denser over time. I will come back to this interesting observation in Section 3.6.

Last, Appendix B.4 shows average wages of entrants to local labor markets relative to the wages of incumbents. This exercise is informative about the pattern of market entrants’ quality across space. An entrant can be anybody who is newly observed in the local labor market in the current period, joining the labor force for the first time, switching from a different region, or entering from unemployment or outside of the labor force. Clearly, under that definition, all relative entry wages are well below zero. This is not surprising because entrants typically have less experience than incumbents and their matching might be worse. What is more surprising, however, is the positive relationship between entry wages and employment growth: the more employment within a local labor market is growing, the higher are relative entry wages. This finding is consistent with the hypothesis that entrants to regions with rising employment are of comparably high quality which is reflected in their comparably high wages.

⁸ Appendix B.4 further shows that the positive relationship still holds when controlling for age, education or occupation.

Nevertheless, the positive correlation does not appear when defining an entrant as a mover joining the region from another local labor market only. Under that definition, the positive correlation falls close to, or even below zero.⁹ In fact, the correlation between employment growth and relative entry wages is mainly driven by labor market entrants. One reason for that might be that schooling systems of growing regions in Southern Germany are comparably better (Combes, Duranton, Gobillon, and Rouxe, 2012).¹⁰ Another explanation could be that most of the sorting of workers happens right after their graduation from school or university.

In summary, with the exception of changes in education shares, regions with

⁸Böhm et al. (2019) find the exact opposite with respect to occupational employment growth: the more an occupation is growing, the lower are the wages of entrants to that occupation relative to the wages of incumbents.

⁹Notice that this might also provide evidence against the hypothesis that growing regions need to pay higher wages to attract workers because of search frictions. If that were true, relative entry wages of movers should also be positively correlated with local employment growth.

¹⁰Consistently, the ratio of students per school is lower in growing regions. I will come back to this later.

growing employment seem to have experienced comparably positive changes in their occupation and industry structure as well as seem to have attracted entrants of relatively high quality. This already suggests that worker quality might have improved in regions with growing employment. The next sections will investigate this in detail.

3.4 Estimating Local Wage Premia and Amenities

This section shows how to estimate changes in location specific wage premia paid for a unit of skill as well as changes in local amenities within a generalized Roy model setting (Roy, 1951; Dahl, 2002). The method applies an insight from Böhm et al. (2019) on the relation between wage growth of individual workers and their location choices in panel data allowing for worker self selection because of changing wage premia and changing amenities at the same time. The method leads to an estimation strategy which is feasible also for a large set of local labor markets.

3.4.1 Model

Worker i at time t faces the problem of choosing in what potential local labor market $l = 1, \dots, L$ to work in. The worker is assumed to choose the region in which his (log) utility would be highest:

$$u_{i,t} = \max\{u_{1,i,t}, \dots, u_{L,i,t}\} = \sum_{l=1}^L I_{l,i,t} u_{l,i,t} \quad (3.1)$$

The realized utility of a worker depends on the chosen location $I_{l,i,t} \equiv \mathbb{1}[u_{l,i,t} \geq u_{j,i,t} \forall j \neq l]$. This choice, in turn, hinges on latent (real) wages $w_{l,i,t}$ the worker can potentially earn in the different regions as well as non monetary local amenities $v_{l,i,t}$ in a linear and additive way.¹¹ This rationalizes that not all workers move to the best paying regions immediately but weigh up wages and amenities (Rosen, 1986; Roback, 1988):

$$u_{l,i,t} = w_{l,i,t} + v_{l,i,t} \quad (3.2)$$

Potential wages $w_{l,i,t}$ depend on the worker's individual skill $s_{i,t}$ common across regions and their contemporaneous return in the local market $\pi_{l,t}$. Hence, for a given unit of skill, the worker might obtain a very different wage premium depending on where he chooses to work. Variation in wage premia might exist in spatial equilibrium because, for instance, demand in some places might be

¹¹Note that log additivity implies that latent utilities in non-log terms are defined as: $U_{l,i,t} = W_{l,i,t}V_{l,i,t}$. This implies that wages and amenities are complements for a worker so that an amenity is valued the more, the higher the wage is. This might reflect that workers can only enjoy amenities if their wage is sufficiently high.

higher than in other places (Bound and Holzer, 2000). In addition, potential real wages depend on a time constant price index r_l which describes differences in price levels across areas.¹² The choice model therefore reflects workers balancing wages, amenities and local prices in equilibrium as is common in the urban economics literature (Glaeser, 2008).

Furthermore, latent amenities hinge on the (overall) local amenity value of a region $\psi_{l,t}$ as well as an idiosyncratic evaluation of a region $\varepsilon_{l,i,t}^v$. This specification rationalizes, for instance, workers switching out of regions with highest wage premia and highest amenities and into low return, low amenity regions for idiosyncratic reasons:

$$w_{l,i,t} = \pi_{l,t} + s_{i,t} - r_l \quad (3.3)$$

$$v_{l,i,t} = \psi_{l,t} + \varepsilon_{l,i,t}^v \quad (3.4)$$

The main aim of this paper is to estimate how local wage premia $\pi_{l,t}$ have evolved over time across different local labor markets as well as how amenities $\psi_{l,t}$ have changed across regions. Böhm et al. (2019) show how one can estimate these parameters in panel data with information on workers' individual wage growth as well their decision in what market to work in.¹³ Building upon their insight and letting Δ denote changes between periods $t-1$ and t , workers' real wage growth can (approximately) be described by a linear equation:

$$\Delta w_{i,t} = \sum_{l=1}^L \bar{I}_{l,i,t} \Delta w_{l,i,t} - \sum_{l=1}^L \Delta I_{l,i,t} \bar{v}_{l,i,t} \quad (3.5)$$

$\bar{I}_{l,i,t} = \frac{I_{l,i,t-1} + I_{l,i,t}}{2}$ is an "average" choice indicator equal to 0 when worker i did neither work in l at t or $t-1$; equal to 1 if he decided to stay in l between these periods; and equal to 0.5 if he either entered or left l at t or $t-1$. Hence, $\bar{I}_{l,i,t}$ describes the average sorting between two periods.¹⁴

In turn, $\Delta I_{l,i,t} = I_{l,i,t} - I_{l,i,t-1}$ contains information on the changing sorting. $\Delta I_{l,i,t}$ equals +1 if i switched into l at t , equals -1 if i left l at t , and is zero otherwise. Because of the general difference between $\Delta I_{l,i,t}$ and $\bar{I}_{l,i,t}$ when there is enough regional mobility, average amenities and changes in potential wages can be separately identified from workers' wage growth and their location decisions relative to a base period $t = T_{base}$ and relative to a base region $l = L_{base}$.

¹²Data on local price differences comes from Kawka et al. (2009) who compute consumer prices by district with prices collected between 2006 and 2008. Unfortunately, up to now, there is no comprehensive panel data available with information on how prices have changed over time and across German regions.

¹³In their application, a market refers to an occupation k . See their Appendix B for the details on the derivation of an estimable equation with respect to both changing occupational skill prices and amenities.

¹⁴Likewise, $\bar{v}_{l,i,t} = \frac{v_{l,i,t} + v_{l,i,t-1}}{2}$ denotes the average amenity value of region l between $t-1$ and t including the systematic part $\bar{\psi}_{l,t} = \frac{\psi_{l,t} + \psi_{l,t-1}}{2}$ as well as the idiosyncratic part $\bar{\varepsilon}_{l,i,t}^v = \frac{\varepsilon_{l,i,t}^v + \varepsilon_{l,i,t-1}^v}{2}$.

Intuition In total, Equation (3.5) consists of two sums. The first summand on the right hand side of Equation (3.5) is a purely pecuniary part: if a worker stays in local labor market l ($\bar{I}_{l,i,t} = 1$, $\Delta I_{l,i,t} = 0$), his observed wage gain equals the potential wage change $\Delta w_{l,i,t}$ in that local labor market. In turn, the potential wage change depends on how market wage premia evolve and how the worker's skill changes.

In contrast, if the worker switches from market l to m ($\bar{I}_{l,i,t} = \frac{1}{2}$, $\Delta I_{l,i,t} = -1$ as well as $\bar{I}_{m,i,t} = \frac{1}{2}$, $\Delta I_{m,i,t} = 1$), one half of his wage growth is approximated to stem from the origin region's potential wage change; and one half of the wage growth to stem from the destination region's potential wage change. The exact allocation of wage growth into origin and destination region is unknown as the data does not contain any information at what exact wage the worker would be indifferent (Böhm et al., 2019).

However, the observed wage change of a moving worker will not only depend on pecuniary motives. Instead, the observed wage change also depends on the amenity differences between regions l and m as amenities serve as compensating differentials in this framework (Rosen, 1986; Roback, 1988). If, for instance, amenity values in Northern Germany are comparably low, wage growth of migrants from Southern to Northern Germany should be relatively high to compensate workers for the potential decline in utility due to low amenities in the destination region. Not accounting for the importance of amenities in location decisions (as is also the case in simple fixed effects models) might therefore severely bias the results as high wage growth of movers would then be attributed to high skill growth.¹⁵

The second summand on the right of Equation (3.5) incorporates the influence of amenities on workers' observed wage growth: with optimal choices, a worker's observed wage growth is the change in the potential wage of the destination region minus the utility gain (or loss) from the behavioral response of switching region. That is, if a utility-optimizing worker chooses to move from the less desirable region l to the more desirable region m , in the data that worker's wage growth should be *lower* than predicted by the changing relevant potential wage alone. Vice versa, observed wage growth should be *higher* than the potential wage change when a worker moves from a more desirable region to a less desirable one.

Average amenity values can therefore be identified by comparing the (conditional) wage growth of movers to the wage growth of stayers. However, I can only identify them relative to a base region which I choose to be the rural region $L_{base} = \text{Husum}$ located in the most northern part of Germany. Mechanically, the reason is that the $\Delta I_{l,i,t}$ sum to zero over all l which leads to multicollinear-

¹⁵Notice that in a fixed effect model, the estimated fixed effect of a worker moving from a high amenity to a low amenity region would be higher compared to a scenario in which amenities play no role. The reason for this is that the worker would be compensated for this less favorable move by a rising wage (Sorkin, 2018). In fact, part of mover's fixed effect would reflect amenity differences between regions instead of worker's time constant ability.

ity. The economic reason is that choices and wages can only be used to identify relative utilities but not utility levels.

Disentangling growth in wage premia from growth in individual worker quality requires a further parameterization of the skill accumulation function:

$$\Delta s_{i,t} = f(X_{i,t-1}) + \varepsilon_{i,t}^s \quad (3.6)$$

$$= X_{i,t-1}\Gamma_{l',l} + \varepsilon_{i,t}^s \quad (3.7)$$

According to Equation (3.7), workers' skill accumulation linearly depends on observables $X_{i,t-1}$ and an unobservable skill shock $\varepsilon_{i,t}^s$ which influences workers' productivity commonly across all regions. I allow skills to change over the life cycle with the speed of skill growth varying by age group (in the empirics, I use dummies for ages 25 – 34, 35 – 44 and 45 – 54). Further, I allow age profiles to vary with previous (l') and current (l) work location. This reflects the fact that skill accumulation might be higher in urban places (Glaeser, 1999; De la Roca and Puga, 2017). As described later, using such a fully saturated skill model also helps in dealing with endogeneity of the choice with respect to error realizations. Inserting this yields the following estimation equation with parameters $\Delta\pi_{l,t}$, $\bar{\psi}_{l,t}$, $\Gamma_{l',l}$ estimable by ordinary least squares.¹⁶

$$\Delta w_{i,t} = \sum_{l=1}^L \bar{I}_{l,i,t} (\Delta\pi_{l,t} + \Delta s_{i,t}) - \sum_{l=1}^L \Delta I_{l,i,t} \bar{v}_{l,i,t} \quad (3.8)$$

$$= \sum_{l=1}^L \bar{I}_{l,i,t} (\Delta\pi_{l,t} + X'_{i,t-1} \Gamma_{l',l}) - \sum_{l=1}^L \Delta I_{l,i,t} \bar{\psi}_{l,t} + \eta_{i,t} \quad (3.9)$$

3.4.2 Identifying Assumptions

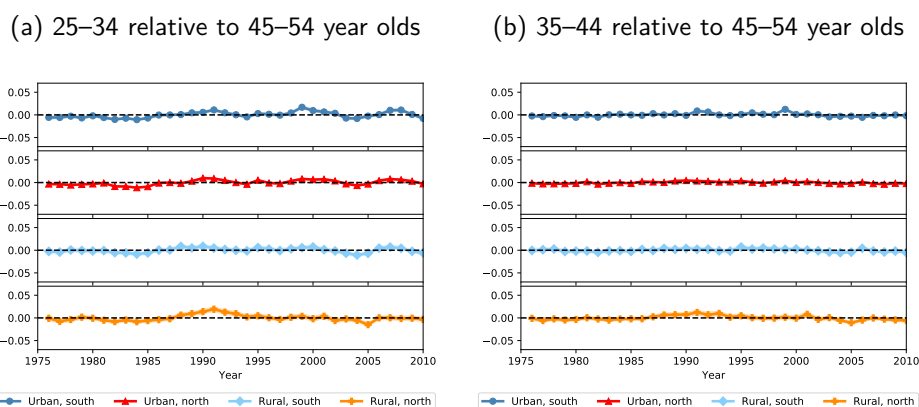
I now make the assumptions explicit which are necessary for a consistent estimation of the main parameters of interest $\Delta\pi_{l,t}$, $\bar{\psi}_{l,t}$ and provide evidence for their plausibility.

Constant Skill Accumulation Function As already noted in Equations (3.6) and (3.7), the function describing the accumulation of workers' skills $\Delta s_{i,t} = f(X_{i,t-1}) + \varepsilon_{i,t}^s = X_{i,t-1}\Gamma_{l',l} + \varepsilon_{i,t}^s$ must not contain a time t subscript to avoid perfectly collinearity between changes in wage premia and skill changes since $\bar{I}_{l,i,t}$ is interacted with both $\Delta\pi_{l,t}$ (i.e., year dummies) as well as $X_{i,t-1}\Gamma_{l',l}$ (i.e., dummies for age groups, previous and current location). This, in turn, implies the assumption that skill growth over the life cycle for a worker starting in a

¹⁶Notice once more that the local price index used to deflate wages across space is time constant because of lacking panel data (Kawka et al., 2009). As the estimation method will exploit changes over time, the price index cancels out in $\Delta w_{i,t}$ for regional stayers, therefore. Wage growth of movers, which is especially relevant for the estimation of amenities, depends on differences in local prices and so does not cancel out, however.

market in 1975 is the same as it was in 2010.¹⁷ Note that this assumption does not make any restrictions on the skills of entering workers so that the average skill can change freely. Unfortunately, this identifying assumption is not testable directly.

Figure 3.3: Wage growth differences



Notes: The lines show average individual wage growth from $t - 1$ to t by year of 25–34 (Figure 3.3a) and 35–44 (Figure 3.3b) year olds minus average wage growth of 45–54 year olds. Results are centered at zero to show trends over time. The shaded areas around the four lines are 95% confidence intervals. The four groups show a classification of local labor markets into northern and southern, rural and urban areas (see Appendix Figure 3.10). Southern areas are located in Baden-Wuerttemberg and Bavaria. The classification of local labor markets into rural and urban is based on BBSR (2014).

Nevertheless, Figure 3.3 shows that wage growth rates between age groups did not change too much during 1985–2010, at least within the four aggregated local labor markets. The figures present a difference-in-difference result showing residual wage growth of young workers (25–34 year olds, Figure 3.3a) and middle aged workers (35–44 year olds, Figure 3.3b) relative to older workers (45–54 year olds) over time. Holding everything else constant in the model, wage growth differences between two samples in different age groups but the same local market should only reflect changes in skill accumulation between these two groups over time as wage premia and common amenities cancel out for workers within a region. The missing of a clear trend in these figures shows that wage growth between age groups (and thereby skill growth rates in the model) did not change much over time.

With respect to the detailed local labor markets, Appendix Figure 3.19 shows a difference-in-difference-in-difference (i.e., triple difference) result. First, I split the sample in the middle (1993). Then I calculate workers' wage growth between two periods and subtract the wage growth of 45–54 year olds from the wage growth of 25–34 year olds and 34–45 year old workers. After that, I take averages before and after 1993 and subtract them from each other. Figure 3.19 shows

¹⁷This assumption is the same as in Acemoglu and Autor (2011) and Yamaguchi (2018). Complementary to that, Heckman et al. (1998) propose to use old workers for estimation of skill price changes as their wage growth might be less confounded by skill growth.

the results of this exercise plotted against employment changes. Admittedly, there is more variation than for the four broad regions so that the assumption of a time constant skill accumulation function is possibly not exactly fulfilled for every local labor market. However, most of the accelerations or decelerations in wage growth after 1993 are very close to zero. Importantly, there seems to be no relation to employment growth.

Base Period With Constant Wage Premia To separately identify changes in wage premia and skill accumulation parameters from each other, a further restriction on the change in wage premia is necessary: there needs to be a period in which wage premia have been constant. Setting $\Delta\pi_{l,t} = 0, t = 1975, \dots, T_{base}$ then allows me to estimate $\Gamma_{l,l}$ in that period. As $\Gamma_{l,l}$ does not change over time by assumption, the conditional wage growth after T_{base} is, hence, informative on the change in wage premia. Analog to Böhm et al. (2019), I set the end of the base period to $T_{base} = 1985$ for having enough variation to estimate the skill accumulation parameters. Controlling for how wage growth evolves over the life cycle, the excess wage growth in $t > T_{base}$ therefore informs the estimates about the changing returns.

Again, this identifying assumption is not testable directly. However, Appendix Figure 3.20 shows that wage and employment changes were much less pronounced between 1975 and 1984. In addition, the figure shows that there has been no change at all in the density size premium taken place between 1975 and 1985.¹⁸ Overall, these results suggest that the time period 1975–1984 was at least a period with less wage and employment changes compared to 1985–2010. In addition to that, Böhm et al. (2019) show that, even when the assumption of constant premia in the base period were violated, the method still identifies an important parameter: namely, how wage premia changed relative to the change that occurred during 1975–1984. That is, how much the change in wage premia accelerated or decelerated over time.

Distribution of Shocks If $\bar{\varepsilon}_{l,i,t}^v$ is not common across all markets $l = 1, \dots, L$ and therefore the locational choice is influenced by the the error draw, a correlation between the choice variables and the error appears which leads to an endogeneity bias because:

$$\eta_{i,t} = \varepsilon_{i,t}^s - \sum_{l=1}^L \Delta I_{l,i,t} \bar{\varepsilon}_{l,i,t}^v \quad (3.10)$$

Nevertheless, l specificity of $\bar{\varepsilon}_{l,i,t}^v$ makes the model much more realistic as workers move in all kinds of directions during their life cycle for presumably idiosyn-

¹⁸I measure this premium as the elasticity of average wages with respect to market density. Market density is defined as the number of employed workers in a local labor market divided by the area of a market in square kilometers. The elasticity was 6.9 log points in 1975 and 7.0 log points in 1985.

cratic reasons.¹⁹ This is why Böhm et al. (2019) show a large set of Monte Carlo experiments which suggest that the bias is of small magnitude in realistic scenarios. The reason for that is that the interaction between previous and current work location in the fully saturated skill accumulation model already picks up most of the endogeneity as the corresponding parameters in $\hat{\Gamma}_{l',l}, l' \neq l$ are identified from the (time constant) sorting of regional movers together with their (biased) wage growth stemming mainly from the error realizations.

In fact, region specificity of $\bar{\varepsilon}_{l,i,t}^v$ introduces a substantial bias in the estimates for the skill accumulation parameters $\hat{\Gamma}_{l',l}, l' \neq l$. However, if the shock distribution stayed constant during the base period and afterwards, the overall bias for the (time varying) wage premium estimates and the (time varying) amenity estimates is minimal: there should be no excess wage growth due to the shocks after the base period compared to the wage growth which took place during the base period (provided the shock distribution remained constant). Notice that this is different for fixed effects models as in Combes, Duranton, and Gobillon (2008) who rely on the exogenous mobility assumption which postulates that workers must not move because of idiosyncratic shock realizations.

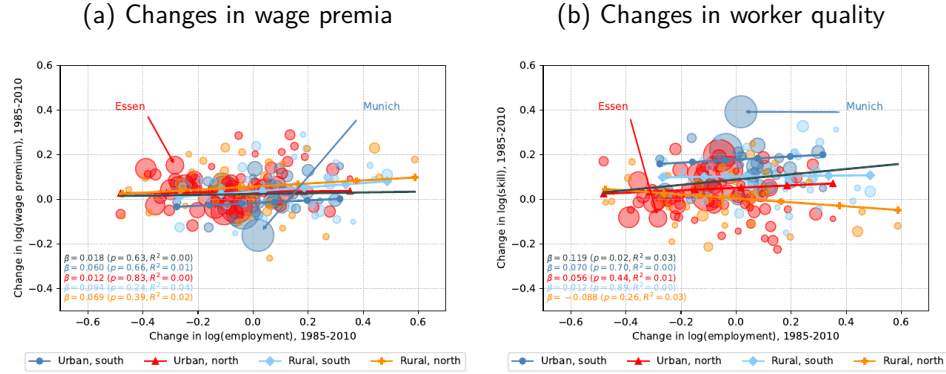
3.5 Local Wage Premia and Worker Quality

This section first presents the estimates of changes in wage premia paid for a unit of skill and changes in worker quality. After that, I show observable determinants of changing worker quality which includes evidence on spatial differences in the occupation, industry and education structure as well as variation in local amenities.

Figure 3.4 shows the main result of this paper. Changes in local wage premia and employment growth between 1985 and 2010 were essentially uncorrelated across local labor markets (p -value = 0.63). Put differently, changes in wage premia paid for a unit of skill were roughly equal between declining and growing regions. This finding is in contrast to Figure 3.2 which showed that growing employment came in hand with growing wages. Figure 3.4b shows the reason for this. It plots the change in worker quality between 1985 and 2010 by market

¹⁹I assume that skill shocks $\varepsilon_{i,t}^s$ have the same influence across all local labor markets. This is in line with the assumption that worker productivity is the same across all regions. Hence, there is no endogeneity problem with respect to the skill shocks as they do not influence workers' location choices.

Figure 3.4: Worker quality increased in growing regions



Notes: Figure 3.4a shows estimated changes in local wage premia between 1985 and 2010. OLS estimates as described by Equation (3.9). Figure 3.4b shows estimated changes in worker quality as described by Equation (3.11). The horizontal axes depict the change in log employment between 1985 and 2010. One bubble represents one of 182 West German local labor markets aggregated from political districts as suggested by Kropp and Schwengler (2011). Bubble size corresponds to the number of employed workers in a market averaged across years 1985 to 2010. The four groups show a classification of local labor markets into northern and southern, rural and urban areas (see Appendix Figure 3.10). Southern areas are located in Baden-Wuerttemberg and Bavaria. The classification of local labor markets into rural and urban is based on BBSR (2014). Regression lines across all areas (black) and within the four broad groups (colored) are weighted by the number of employed workers.

as the residual between changes in wage premia and changes in average wages:

$$\sum_{t=1986}^{2010} \underbrace{E[s_{l,i,t}|I_{l,i,t}=1] - E[s_{l,i,t-1}|I_{l,i,t-1}=1]}_{\text{Mean skill change in } l} = \quad (3.11)$$

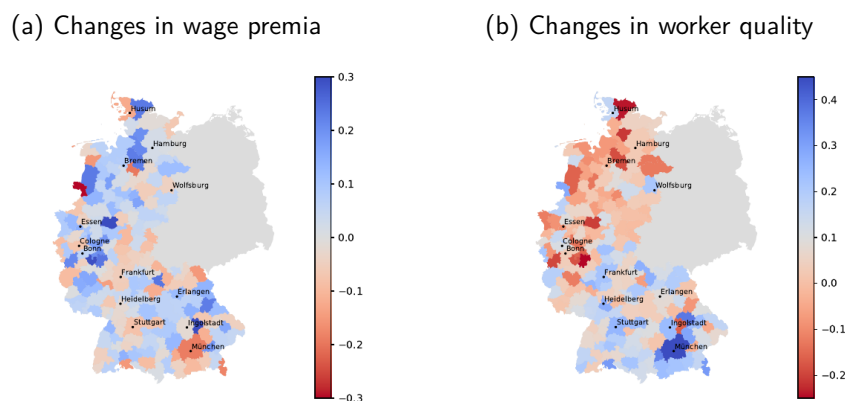
$$\sum_{t=1986}^{2010} \underbrace{E[w_{i,t}|I_{l,i,t}=1] - E[w_{i,t-1}|I_{l,i,t-1}=1]}_{\text{Mean wage change in } l}$$

$$- \sum_{t=1986}^{2010} \underbrace{\Delta\pi_{l,t}}_{\text{Change in wage premium}}$$

Worker quality strongly increased in regions with growing employment. On average, a market whose employment increased by 10% experienced an increase in worker quality of 1.2%. This finding is in line with the hypothesis that growing regions were able to attract workers of comparably high quality as already suggested in Section 3.3.2 that showed remarkable differences in changes of the occupation and industry structure between growing and shrinking regions.

Figure 3.5 shows the differential changes in wages premia and worker quality across Germany. There appears to be no clear pattern with respect to changes

Figure 3.5: Geography of changes in wage premia and skills, 1985 – 2010



Notes: Figure 3.5a shows estimated changes in local wage premia between 1985 and 2010. OLS estimates as described by Equation (3.9). Figure 3.5b shows estimated changes in worker quality as described by Equation (3.11). Assignment of political districts to West German local labor markets is based on commuting flows and follows Kropp and Schwengler (2011). East Germany is left out of the analysis.

in wage premia (Figure 3.5a). In contrast, Figure 3.5b suggests that worker quality primarily increased in Southern Germany.²⁰

The general pattern of improving worker quality also appears when using different estimation methods. Appendix D shows the results of four robustness checks. First, I drop the terms in Equation (3.9) which control for changing local amenities. This increases the elasticity between skill and employment changes further. Second and third, I allow for a more flexible skill accumulation function by including dummies for occupation and education separately. This also slightly increases the elasticity between worker quality changes and employment changes. Last, I estimate changes in wage premia using a fixed effects approach following Cortes (2016). I identify changes in wage premia from year times region dummies and control for unobserved worker quality with worker times region fixed effects. Additionally, I allow for region specific, concave wage profiles over the life cycle by adding age times region dummies to the regression. Using this complementary approach, the estimated skill elasticity rises further.

In summary, in all of these robustness checks, there is a zero (or even negative) correlation between changing wage premia and employment growth. In contrast, employment growth and skill changes are strongly correlated in all specifications.

²⁰Appendix Figure 3.21 shows the accumulated skill of a hypothetical stayer by market density. In line with De la Roca and Puga (2017), I find that workers in urban places have higher skill accumulation rates. However, there is also a lot of variation between urban places. For instance, stayers' estimated skills at age 54 in Munich are roughly 60 log points higher than in Essen. Part of this difference might be explained by differences in the occupation structure with high shares of low accumulation, production jobs in Essen and high accumulation, professional jobs in Munich. See Böhm et al. (2019) for skill profiles by occupation.

There are several (model independent) explanations for the positive correlation between worker quality changes and employment growth as well as the concentration of positive worker quality changes in Southern Germany: shifts in observable characteristics such as the occupation, education and industry structure might have been different between booming and declining regions (Findeisen and Südekum, 2008); increasing density of employment could have raised agglomeration economies in booming regions making workers and firms become more productive (Duranton and Puga, 2004); there might be an important role which local amenities play to attract skilled workers (Glaeser, Kolko, et al., 2001; Diamond, 2016).

My aim in this paper is not to fully disentangle these explanations or even uncover causal relations that determine the performance of local labor markets. Instead, I will develop new stylized facts about the processes of changing worker quality and employment growth: what have been the main (observable) characteristics of successful places that experienced a growth in all three: wages, worker quality and employment? I will focus on the role played by occupations, industries and education as well as local amenities in greater depth.

3.5.1 The Role of a Shifting Occupation, Industry, and Education Structure

It is well documented across a wide range of countries that the distribution of occupational and industrial employment has shifted a lot of over the last few decades. Employment declined in routine, producing occupations and increased in analytical, professional as well as manual, service jobs (Dustmann, Ludsteck, et al., 2009). Similarly, employment decreased in manufacturing industries and rose in high- as well as low-skill services (Bárány and Siegel, 2018).

What has been less well documented is the fact that changes in occupational and industrial employment are far from evenly distributed across regions within countries (Autor, 2019). Section 3.3.2 already showed up huge differences in changes of the occupation and industry structure across space. For instance, although the countrywide share of employment within producing occupations decreased from 58 percentage points to 50 points over time in Germany, the shifts across local labor markets range from -18 to +3 percentage points (see Appendix Figure 3.11c). In this section, I will now investigate the impact of differential changes in the occupation, industry, and education structure on changes in worker quality.

Panel A in Table 3.1 shows the results from a regression of the estimated changes in worker quality between 1985 and 2010 on changes in occupational (log) employment. Variation comes from differential changes in occupational employment as well as quality changes between local labor markets. Worker quality increased in regions which experienced an increase in managerial, professional and technical (Mgr-Prof-Tech) employment. In contrast, worker quality declined where low wage, service and care (Srv-Care) employment rose as well

Table 3.1: Observable determinants of changing worker quality

		$\Delta \bar{s}_{i,2010}$	$\Delta \bar{s}_{i,2010}$	$\Delta \bar{s}_{i,2010}$	$\Delta \bar{s}_{i,2010}$	$\Delta \bar{s}_{i,2010}$	$\Delta \bar{s}_{i,2010}$	$\Delta \bar{s}_{i,2010}$
<i>Panel A</i>								
Occupation growth	Mgr-Prof-Tech	0.30 (0.00)						0.14 (0.04)
	Sales-Office	0.08 (0.06)						-0.07 (0.26)
	Prod-Op-Crafts	-0.36 (0.00)						-0.52 (0.01)
	Srvc-Care	-0.04 (0.10)						-0.09 (0.00)
<i>Panel B</i>								
Industry growth	Machine building		0.06 (0.01)					0.08 (0.02)
	High-skill services		0.19 (0.00)					0.21 (0.00)
	Manufacturing		-0.15 (0.00)					0.04 (0.31)
	Low-skill services		-0.07 (0.11)					0.07 (0.10)
<i>Panel C</i>								
Education growth	University degree			0.13 (0.00)				-0.09 (0.04)
	Medium educated			-0.15 (0.01)				0.12 (0.61)
	Low educated			0.03 (0.16)				0.00 (0.86)
<i>Panel D</i>								
Amenities	Estimated $\bar{\psi}_{i,2010}$				2.23 (0.00)			
	Amenity index					0.68 (0.00)		0.65 (0.00)
	Criminal offenses						-0.08 (0.02)	0.01 (0.72)
	Students per school						0.07 (0.08)	0.04 (0.31)
Adjusted R^2		0.35	0.20	0.06	0.13	0.33	0.03	0.01
N		182	182	182	182	182	182	182


Notes: The panels show results from regressions of estimated worker quality changes on variables varying across panels. Variation comes from differential changes across local labor markets between 1985 and 2010. Changes in worker quality are estimated according to Equation (3.11). Results are weighted by the number of employed workers in a market averaged across years 1985 to 2010. p -values are in parentheses. Occupations: Mgr-Prof-Tech: managers, professionals, technicians; Prod-Op-Crafts: production, operators, craftsmen; Srvc-Care: services and care. Classification of detailed occupations follows Acemoglu and Autor (2011), see Böhm et al. (2019) for the exact mapping. Industries: Machine building: manufacture of basic metals and fabricated metal products, machinery and equipment, electrical and optical equipment, transport equipment; High skill services: education, health and social, public administration and defense, compulsory social security, financial intermediation, real estate, renting and business activities, electricity, gas and water supply, other community, social and personal service activities, communication; Manufacturing: manufacture of wood and wood products, coke, refined petroleum products and nuclear fuel, chemicals; Low-skill services: hotels, restaurants, wholesale, food. Education: University degree: technical college or university; Medium educated: Abitur or apprenticeship training; Low educated: without postsecondary education or missing. Amenities: $\bar{\psi}_{i,2010}$ are the estimated amenity values according to Equation (3.9). The amenity index was computed by prognos (2018) combining information on 53 amenity values ranging from sunshine duration to restaurant and physician density. Data on criminal offenses comes from Bundeskriminalamt (2018). The measure was computed as offenses divided by population collected from Statistische Ämter des Bundes und der Länder (2017a). Information on students per school was retrieved from Statistische Ämter des Bundes und der Länder (2017b).

as the number of workers employed in production, operator and crafts (Prod-Op-Crafts) occupations increased.

Similar results apply to the influence of changes in the industrial and education structure (Panels B and C). Worker quality surged where the amount of workers employed in high-skill service industries and machine building expanded as well as the number of university graduates increased. Hence, only regions which were able to “reinvent” (Glaeser, 2005) themselves through changing its employment structure where able to attract workers of high quality. These results are very similar to the findings in Findeisen and Südekum (2008).

To quantify the importance of changes in the occupation and industry structure, Figure 3.6 repeats the decomposition of changes in wages into changing wage premia and worker quality according to Equation (3.11). In addition to Figure 3.4, I reweighted observations when computing average wages to eliminate the impact of geographic differences in the number of workers employed

Figure 3.6: Reweighted skill changes



Notes: Figure 3.6a shows estimated changes in worker quality as described by Equation (3.11). Observations were reweighted following DiNardo et al. (1996) to eliminate differential changes in the occupation structure across local labor markets. Figure 3.6b repeats this by reweighting observations to eliminate differential changes in the industry times occupation structure across local labor markets. The horizontal axes depict the change in log employment between 1985 and 2010. One bubble represents one of 182 West German local labor markets aggregated from political districts as suggested by Kropp and Schwengler (2011). Bubble size corresponds to the number of employed workers in a market averaged across years 1985 to 2010. The four groups show a classification of local labor markets into northern and southern, rural and urban areas (see Appendix Figure 3.10). Southern areas are located in Baden-Wuerttemberg and Bavaria. The classification of local labor markets into rural and urban is based on BBSR (2014). Regression lines across all areas (black) and within the four broad groups (colored) are weighted by the number of employed workers.

in certain occupations, as well as industry times occupation cells, following DiNardo et al. (1996).²¹ Compared to the baseline density of skill changes with respect to employment of 12 log points, both elasticities are smaller and become insignificant (p -values rise from 0.02 in the unweighted scenario to 0.07 and 0.24). Reweighting observations with occupation shares of 1985 induces the elasticity to drop to eight log points. Further reweighting observations to match the combined occupation times industry structure of 1985 decreases the elasticity further to five log points. Overall, this shows that roughly one third to one half of the increase in skill within expanding regions is due the disproportionate growth of high paying managerial, professional and technical as well as machine building employment within expanding regions.

3.5.2 The Role of Local Amenities to Attract Skilled Workers

The urban economics literature has always put focus on the role of local amenities for the locational choices of the workforce and the spatial equilibrium

²¹Reweighting observations with education shares of 1985 does not substantially affect the elasticity which decreases to 11 log points, see Appendix Figure 3.26. The reason for this is that education and employment growth are virtually uncorrelated (Appendix Figure 3.15). Appendix Figure 3.26 also shows the results for reweighting with industries alone instead of industry times occupation cells. The elasticity amounts to 7 log points.

(Glaeser and Maré, 2001; Glaeser, 2011). Following this literature, I will investigate the correlation between the estimated amenity changes as well as information on collected (time constant) amenity proxies as is also done by Diamond (2016).

Appendix Figure 3.27 shows the distribution of estimated amenities, an amenity index based on information on 53 categories ranging from sunshine duration to restaurant and physician density, data on crime rates as well as students per school.²² In general, local amenities appear to be higher in Southern than Northern Germany. This is true for both the estimated amenities as well the amenity index. Consistently, crime rates are lower in the south. In addition, the number of students per school is also lower in Southern Germany.

The results from a bivariate regression of estimated quality changes on amenities are depicted in Panel D of Table 3.1. Again, variation comes from differences in average skill changes and amenities across local labor markets. Regions with high estimated amenity values experienced increases in skill. The same holds for the amenity index. Contrary, areas with high crime rates exhibited lower skill growth. The relation to the number of students per school is not significant at the 5% level. This might indicate that differences in schooling quality between regions (e.g., Baumert et al., 2013) play a minor role for the variation in worker quality changes. In total, this exercise shows that regions which experienced an increase in worker quality also had higher local amenity values. An explanation for this could be the increasing sorting of high quality workers into regions with attractive living standards (Diamond, 2016).

Taking all external explanatory variables together, 59% of the variation in worker quality changes can be explained by these observables.

3.6 Sources of the Rising Density Wage Premium


This section analyses the impact of changing regional wage premia and worker quality on the increasing wage disparities between German local labor markets.

There is a sizable wage premium which workers receive for working in more dense areas across countries (Glaeser and Maré, 2001; Combes, Duranton, and Gobillon, 2008; De la Roca and Puga, 2017; Dauth, Findeisen, Moretti, et al., 2019). Germany is no exception. On the contrary, the German density wage premium has even increased over time. Figure 3.7 shows that having worked in a 100 log point denser region (the difference in density between Munich and Göttingen) was associated with seven log points higher annual wages in 1985.²³

²²The amenity index was computed by prognos (2018). Data on crime rates per person is from Bundeskriminalamt (2018). Information on students per school was retrieved from Statistische Ämter des Bundes und der Länder (2017b). Reassuringly, there is a strong positive correlation between estimated amenities $\bar{\psi}_{l,2010}$ and the amenity index (p -value < 0.01 , see Appendix Figure 3.28). This provides additional plausibility for the estimation approach.

²³Density refers to number of employed workers in a region divided by area measured in square kilometers.

Figure 3.7: Density wage premium increased over time



Notes: Figure 3.7a shows average wages in 1985 plotted against market density defined as number of employed workers divided by area in square kilometers. Figure 3.7b repeats this for 2010. One bubble represents one of 182 West German local labor markets aggregated from political districts as suggested by Kropp and Schwengler (2011). Bubble size corresponds to the number of employed workers in a market averaged across years 1985 to 2010. The four groups show a classification of local labor markets into northern and southern, rural and urban areas (see Appendix Figure 3.10). Southern areas are located in Baden-Wuerttemberg and Bavaria. The classification of local labor markets into rural and urban is based on BBSR (2014). Regression lines across all areas (black) and within the four broad groups (colored) are weighted by the number of employed workers.

This elasticity increased until 2010 by roughly 17%. Is the source of that rise an increased sorting of high quality workers to dense places like Essen and Munich? Or did wage premia rise more in dense places because of higher skill demand?

Figure 3.8a shows the change in estimated wage premia plotted against log density in 1985. The elasticity of price changes with respect to log density is negative. A 100 log point denser region experienced a 1.8 log point lower increase in wage premia. One reason for that could be that dense areas have higher estimated amenity values and so firms in dense places do not need to pay higher wage premia to compensate workers (see Appendix Figure 3.29 for the positive correlation between estimated amenities and market density). Instead, workers might sort into dense places because of these higher amenities and give up rising wage premia they could have earned in less dense places (Roback, 1988).

Notice, however, that the pattern of more favorable amenities in denser places does not appear with respect to the externally collected amenity data. This might indicate the necessity of identifying amenities from worker choices within a revealed preference approach instead of relying on external measures.²⁴ In fact, when I estimate wage premia according to Equation (3.9) but omit the controls for amenities, the interpretation turns around. Appendix Figure 3.30

²⁴Alternatively, this might also reflect that amenities increased more in cities and this increase has been relevant for the changing worker sorting (instead of absolute amenity values). This would not be visible (by definition) in the externally collected amenity proxies as they are time constant and refer to the values at the end of the sample period.

Figure 3.8: Changes in wage premia and worker quality by 1985's density

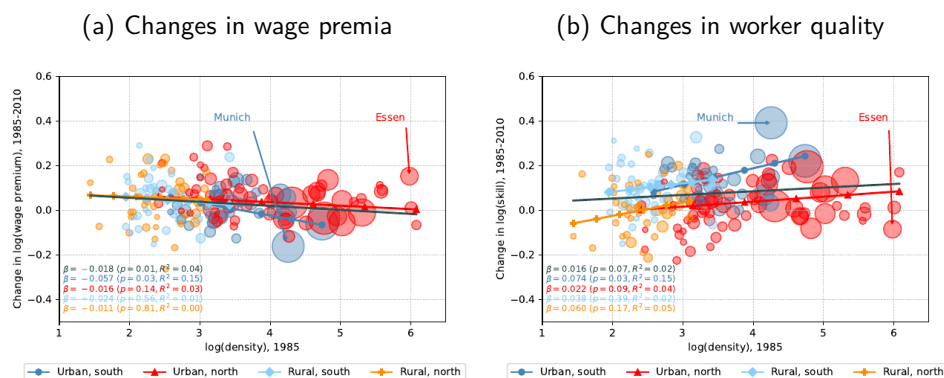


Figure 3.8a shows estimated changes in wage premia between 1985 and 2010. OLS estimates as described by Equation (3.9). Figure 3.8b shows estimated changes in worker quality as described by Equation (3.11). The horizontal axes depicts log employment density 1985 defined as number of employed workers in a market divided by area in square kilometers. One bubble represents one of 182 West German local labor markets aggregated from political districts as suggested by Kropp and Schwengler (2011). Bubble size corresponds to the number of employed workers in a market averaged across years 1985 to 2010. The four groups show a classification of local labor markets into northern and southern, rural and urban areas (see Appendix Figure 3.10). Southern areas are located in Baden-Wuerttemberg and Bavaria. The classification of local labor markets into rural and urban is based on BBSR (2014). Regression lines across all areas (black) and within the four broad groups (colored) are weighted by the number of employed workers.

shows the results: when omitting the terms which control for local amenities, I find that denser places experienced larger increases in wage premia. This also holds when estimating changes in wage premia via year times region fixed effects while controlling for unobserved worker heterogeneity via worker times region fixed effects similar to Combes, Duranton, and Gobillon (2008) or Cortes (2016). This underscores the importance of a method which is able to distinguish worker sorting because of pecuniary motives from worker sorting because of non-pecuniary motives. When ignoring the influence of amenities, estimated changes in wage premia might be biased as workers also sort with respect to non-pecuniary aspects instead of potential wages alone. The sign of the bias, in turn, depends on the correlation between amenities and market density.

In contrast to the negative wage premium elasticity (again, which only appears when controlling for local amenities), the relation between skill changes and density is positive (Figure 3.8b). A 100 log point denser region experienced a 1.6 log point higher increase in worker quality. This effect thereby almost exactly offset the relative decline in wage premia within dense areas.

If changes in wage premia and worker quality balanced each other, why then did the density premium increase over time? The reason for this is that employment in dense German regions did not rise further. On the one hand, previously dense places such as the Ruhrgebiet including Essen lost large fractions of their employment force as shown in Figure 3.9a. On the other hand, many less dense areas – mainly in southern Germany – gained employment and

Figure 3.9: Changes in employment and wages by 1985's density

(a) Density declined in formerly dense areas (b) Wage changes unrelated to 1985's densities




Figure 3.9a shows changes in employment density of a local labor market between 1985 and 2010. Figure 3.9b shows the change in average wages. The horizontal axes depicts log employment density 1985 defined as number of employed workers in a market divided by area in square kilometers. One bubble represents one of 182 West German local labor markets aggregated from political districts as suggested by Kropp and Schwengler (2011). Bubble size corresponds to the number of employed workers in a market averaged across years 1985 to 2010. The four groups show a classification of local labor markets into northern and southern, rural and urban areas (see Appendix Figure 3.10). Southern areas are located in Baden-Wuerttemberg and Bavaria. The classification of local labor markets into rural and urban is based on BBSR (2014). Regression lines across all areas (black) and within the four broad groups (colored) are weighted by the number of employed workers.

so became denser over time. In combination with almost density independent changes in average wages (see Figure 3.9b showing the sum of Figures 3.8a and 3.8b for completeness), this made the slope of the density wage premium become steeper over time. Therefore, the rise in the premium is mainly because of declining employment in the formerly densest areas as well as expanding employment in previously less dense places. Changes in wage premia and changes in worker quality, however, have offset each other on average across markets of different density.

3.7 Conclusion

There has been considerable variation in wage and employment changes across German local labor markets during the last few decades. Whereas southern regions experienced both an increase in wages and employment on average, the opposite was true for many northern regions. This resulted in a strong connection between wage and employment growth across markets, seemingly consistent with geographic differences in shocks to demand. This paper has investigated the drivers of these opposing trends in more detail: did demand increase relatively more in regions with growing employment? Or did worker quality improve in expanding areas?

To separate these two competing explanations, I have estimated selection

corrected local wage premia paid for a unit of skill as well as differences in local amenities within a generalized Roy model setting. In that framework, workers decide in what region to work in based on latent, potential real wages and local amenities. The model was informed by variation in workers' wage growth and their location decisions including workers' moving behavior.

The results have shown that, contrary to average wages, quality adjusted wage premia have not increased more in expanding regions. That is, the increase in wage premia workers get paid for their skills was roughly the same in growing and shrinking local labor markets. Instead, the quality of workers employed in regions with expanding employment improved. This resulted in comparably high wage growth within these regions.

I have investigated several potential explanations for this finding by comparing observables of regions with growing worker quality to the observables of regions with declining skill. Worker quality improved in regions which were able to attract professional and technical occupations as well as firms engaged in the machine building sector. This primarily happened in Southern Germany; a region which was able to "reinvent" itself over time (Findeisen and Südekum, 2008). Average skill also increased in regions with high estimated amenity values. This suggests that skilled workers might sort into regions with attractive living standards although local prices might be higher in these areas as well (Diamond, 2016); a possibility for which I control in the estimation.

Last, I have explored the determinants of the rising market density premium. I found that worker quality has increased more in dense places. This effect, however, was offset by a relative decline in wage premia of dense places. A potential explanation for that might be the rising importance of amenities in the location decisions of workers and their wage compensating nature. In line, dense places have higher estimated amenity values in the estimation. This pattern does not appear with respect to external data on amenity values, however.

Further disentangling possible mechanisms which influence the sorting of skill across space is a fruitful topic for future research and policy. The answer to the question about how local labor markets can attract skilled workers and successful firms is of first order importance in a world of rising disparities (Arntz, 2010; Moretti, 2012); especially in the light of increasing political polarization between regions because of regional differences in the vulnerability to technical change and competition from abroad (Autor et al., 2016).

Additionally, it would be interesting to investigate possible sources for spatial differences in occupational and industrial change. In a first step, one might therefore estimate how the returns to occupational skill have changed over time and across regions. For instance, Manning and Petrongolo (2017) show that local labor markets are indeed quite local and so should be the returns to skill. Holding the composition of skill constant, such an exercise would then also allow to investigate possible sources of changing occupational skill prices by exploiting additional variation across local labor markets. Such regional variation could range from market thickness in the supply of potential skill (Bleakley and Lin,


2012) to technology investments (Czernich et al., 2011).

A Data Appendix

A.1 Aggregation of Local Labor Markets

Assignment of political districts to West German local labor markets is based on commuting flows and follows Kropp and Schwengler (2011). East Germany is left out of the analysis. Local labor markets are further classified into the groups: urban - north (red), rural - north (orange), urban - south (darkblue) and rural - south (lightblue). Southern areas are located in Baden-Wuerttemberg and Bavaria. The classification of local labor markets into rural and urban is based on BBSR (2014). Figure 3.10 shows the classification across Germany.

Figure 3.10: Spatial distribution of rural and urban Regions



A.2 Aggregation of Occupations

Occupations are categorized into:

1. Managers-Professionals-Technicians (Mgr-Prof-Tech)
2. Sales-Office (Sales-Office)
3. Production-Operators-Craftsmen (Prod-Op-Crafts)
4. Services-Care (Srcv-Care)

Classification of detailed occupations follows Acemoglu and Autor (2011), see Böhm et al. (2019) for the exact mapping.

A.3 Aggregation of Industries

Machine building: manufacture of basic metals and fabricated metal products, machinery and equipment, electrical and optical equipment, transport equipment; High skill services: education, health and social, public administration and defense, compulsory social security, financial intermediation, real estate, renting and business activities, electricity, gas and water supply, other community, social and personal service activities, communication; Manufacturing: manufacture of wood and wood products, coke, refined petroleum products and nuclear fuel, chemicals; Low-skill services: hotels, restaurants, wholesale, food. This largely follows Bárány and Siegel (2018) with the additional differentiation between manufacturing and machine building as, compared to other countries, Germany has a very large and increasing share of machine building firms (Dauth, Findeisen, and Südekum, 2017).

A.4 Education

University degree: technical college or university; Medium educated: Abitur or apprenticeship training; Low educated: without postsecondary education or missing. This classification follows Fitzenberger et al. (2006).

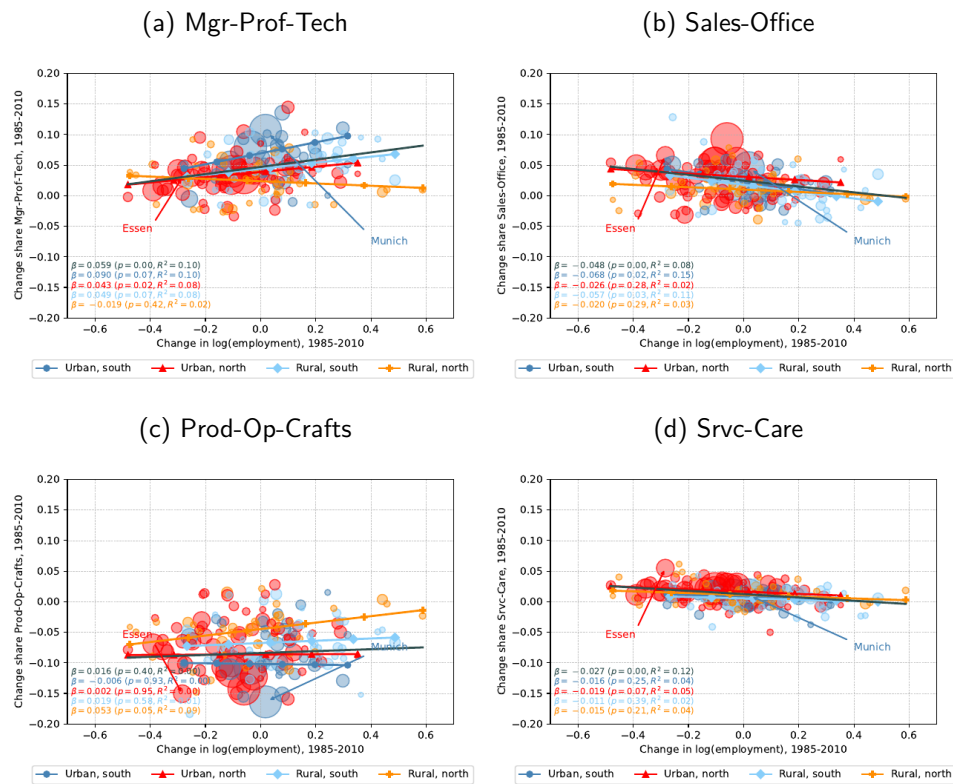
A.5 Amenity Data

The amenity index was computed by prognos (2018) combining information on 53 amenity values ranging from sunshine duration to restaurant and physician density. Data on criminal offenses comes from Bundeskriminalamt (2018). The measure was computed as offenses divided by population collected from Statistische Ämter des Bundes und der Länder (2017a). Information on students per school was retrieved from Statistische Ämter des Bundes und der Länder (2017b). Data on a local price index comes from Kawka et al. (2009). Local price differences correspond to 2006 – 2008.

B Observables of Growing Regions

B.1 Occupation Growth

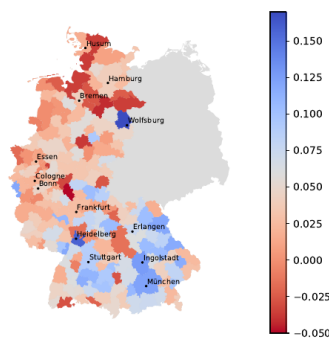
Figure 3.11: Changes of occupational employment



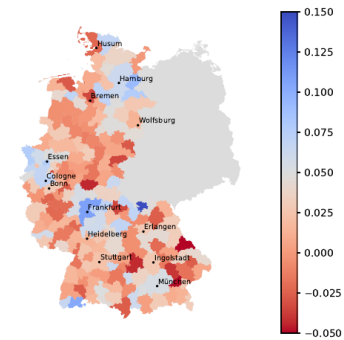
Notes: The vertical axes show changes in occupational employment shares between 1985 and 2010. The horizontal axes depict the change in log employment between 1985 and 2010. Mgr-Prof-Tech: managers, professionals, technicians; Prod-Op-Crafts: production, operators, craftsmen; Srvc-Care: services and care. Classification of detailed occupations follows Acemoglu and Autor (2011), see Böhm et al. (2019) for the exact mapping. One bubble represents one of 182 West German local labor markets aggregated from political districts as suggested by Kropp and Schwengler (2011). Bubble size corresponds to the number of employed workers in a market averaged across years 1985 to 2010. The four groups show a classification of local labor markets into northern and southern, rural and urban areas (see Appendix Figure 3.10). Southern areas are located in Baden-Wuerttemberg and Bavaria. The classification of local labor markets into rural and urban is based on BBSR (2014). Regression lines across all areas (black) and within the four broad groups (colored) are weighted by the number of employed workers.

Figure 3.12: Changes of occupational employment across space

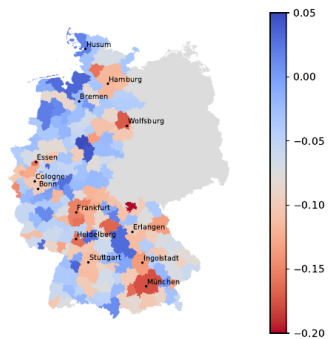
(a) Mgr-Prof-Tech



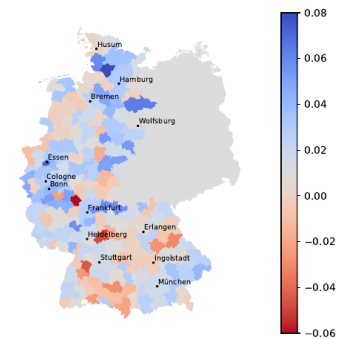
(b) Sales-Office



(c) Prod-Op-Crafts



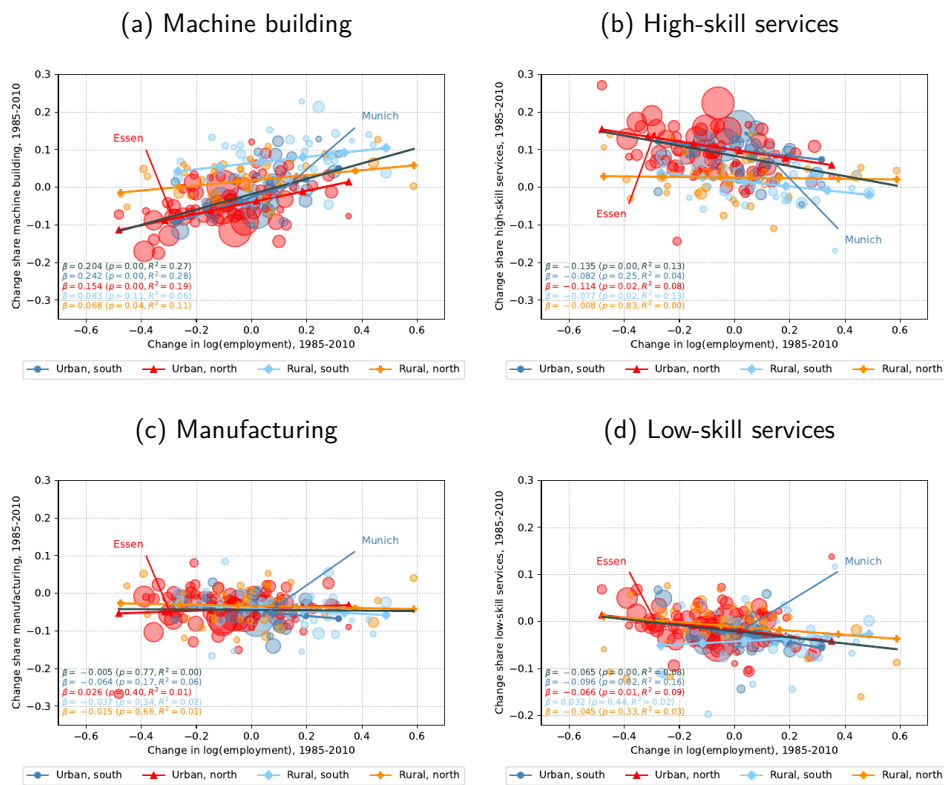
(d) Srv-Care



Notes: The maps show changes in occupational employment shares between 1985 and 2010 across space. Mgr-Prof-Tech: managers, professionals, technicians; Prod-Op-Crafts: production, operators, craftsmen; Srv-Care: services and care. Classification of detailed occupations follows Acemoglu and Autor (2011), see Böhm et al. (2019) for the exact mapping. Assignment of political districts to West German local labor markets is based on commuting flows and follows Kropp and Schwengler (2011). East Germany is left out of the analysis.

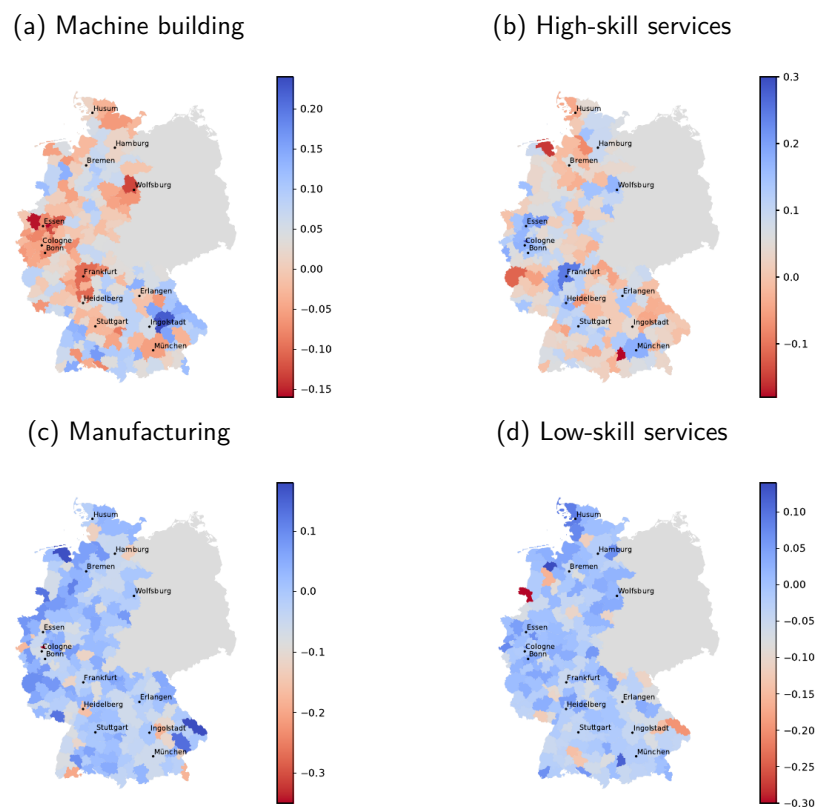
B.2 Industry Growth

Figure 3.13: Changes of industrial employment



Notes: The vertical axes show changes in industrial employment shares between 1985 and 2010. The horizontal axes depict the change in log employment between 1985 and 2010. Machine building: manufacture of basic metals and fabricated metal products, machinery and equipment, electrical and optical equipment, transport equipment; High skill services: education, health and social, public administration and defense, compulsory social security, financial intermediation, real estate, renting and business activities, electricity, gas and water supply, other community, social and personal service activities, communication; Manufacturing: manufacture of wood and wood products, coke, refined petroleum products and nuclear fuel, chemicals; Low-skill services: hotels, restaurants, wholesale, food. One bubble represents one of 182 West German local labor markets aggregated from political districts as suggested by Kropp and Schwengler (2011). Bubble size corresponds to the number of employed workers in a market averaged across years 1985 to 2010. The four groups show a classification of local labor markets into northern and southern, rural and urban areas (see Appendix Figure 3.10). Southern areas are located in Baden-Wuerttemberg and Bavaria. The classification of local labor markets into rural and urban is based on BBSR (2014). Regression lines across all areas (black) and within the four broad groups (colored) are weighted by the number of employed workers.

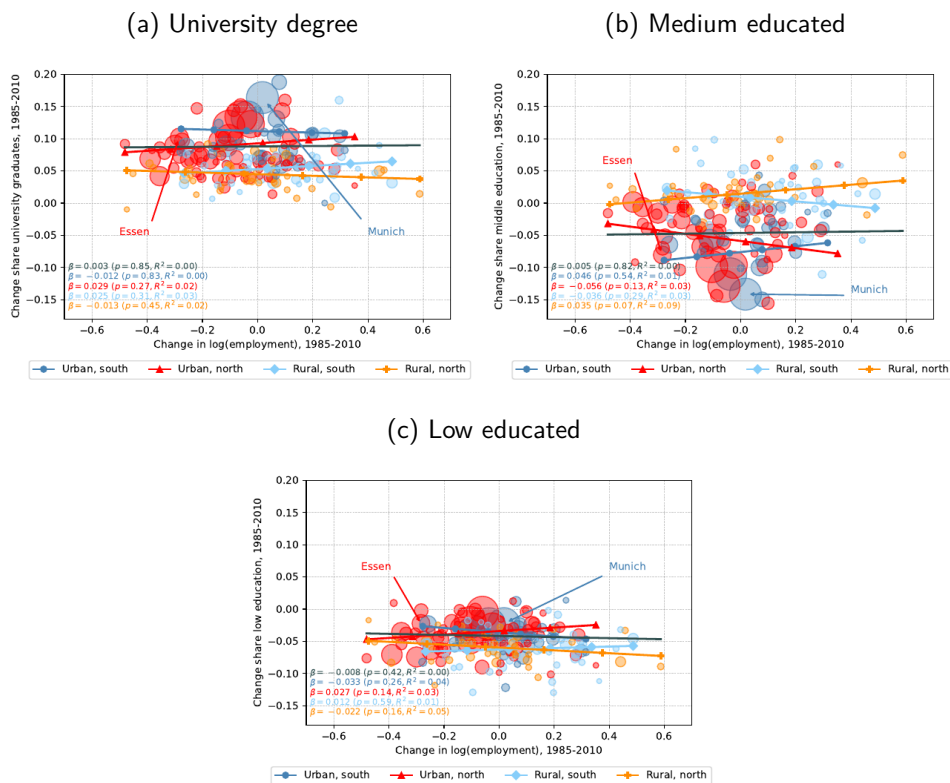
Figure 3.14: Changes of industrial employment across space



Notes: The maps show changes in industrial employment shares between 1985 and 2010 across space. Machine building: manufacture of basic metals and fabricated metal products, machinery and equipment, electrical and optical equipment, transport equipment; High skill services: education, health and social, public administration and defense, compulsory social security, financial intermediation, real estate, renting and business activities, electricity, gas and water supply, other community, social and personal service activities, communication; Manufacturing: manufacture of wood and wood products, coke, refined petroleum products and nuclear fuel, chemicals; Low-skill services: hotels, restaurants, wholesale, food. Assignment of political districts to West German local labor markets is based on commuting flows and follows Kropp and Schwengler (2011). East Germany is left out of the analysis.

B.3 Educational Changes

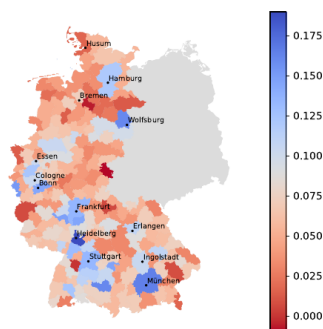
Figure 3.15: Changes of education shares



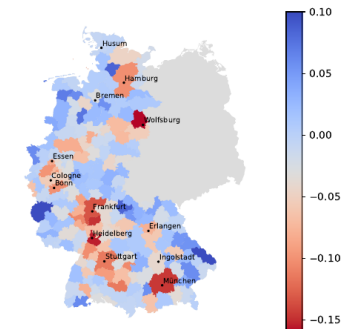
Notes: The vertical axes show changes in education shares between 1985 and 2010. The horizontal axes depict the change in log employment between 1985 and 2010. University degree: technical college or university; Medium educated: Abitur or apprenticeship training; Low educated: without postsecondary education or missing. One bubble represents one of 182 West German local labor markets aggregated from political districts as suggested by Kropp and Schwengler (2011). Bubble size corresponds to the number of employed workers in a market averaged across years 1985 to 2010. The four groups show a classification of local labor markets into northern and southern, rural and urban areas (see Appendix Figure 3.10). Southern areas are located in Baden-Wuerttemberg and Bavaria. The classification of local labor markets into rural and urban is based on BBSR (2014). Regression lines across all areas (black) and within the four broad groups (colored) are weighted by the number of employed workers.

Figure 3.16: Changes of education shares across space

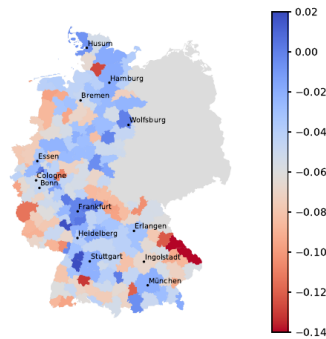
(a) University degree



(b) Medium educated



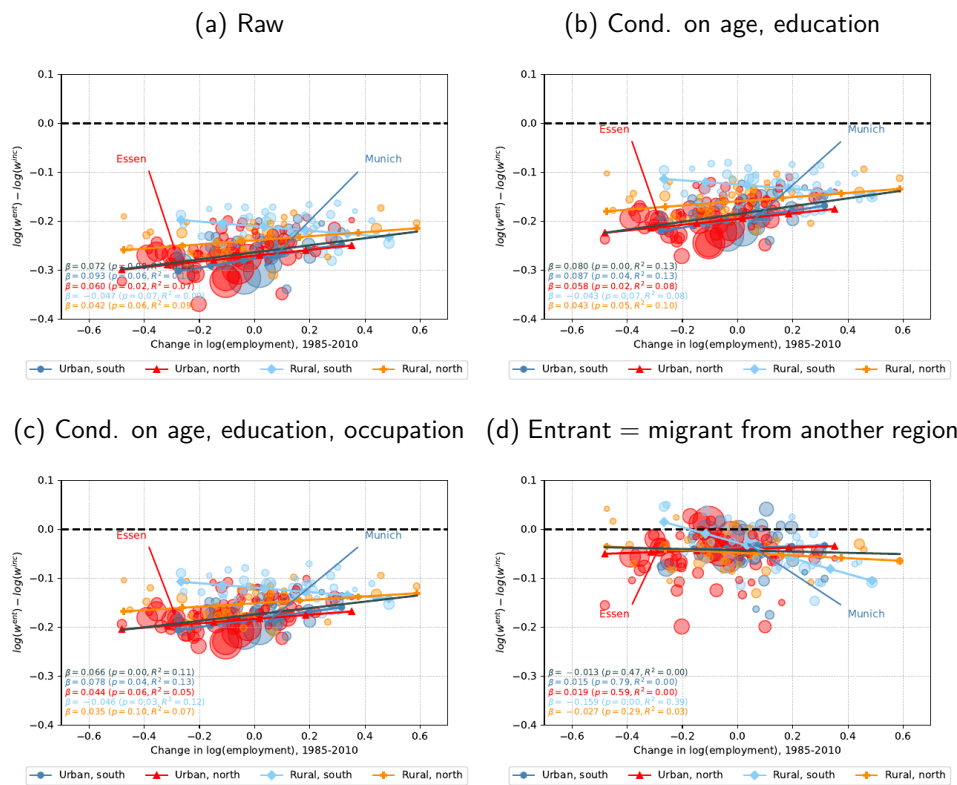
(c) Low educated



Notes: The maps show changes in education shares between 1985 and 2010 across space. University degree: technical college or university; Medium educated: Abitur or apprenticeship training; Low educated: without postsecondary education or missing. Assignment of political districts to West German local labor markets is based on commuting flows and follows Kropp and Schwengler (2011). East Germany is left out of the analysis.

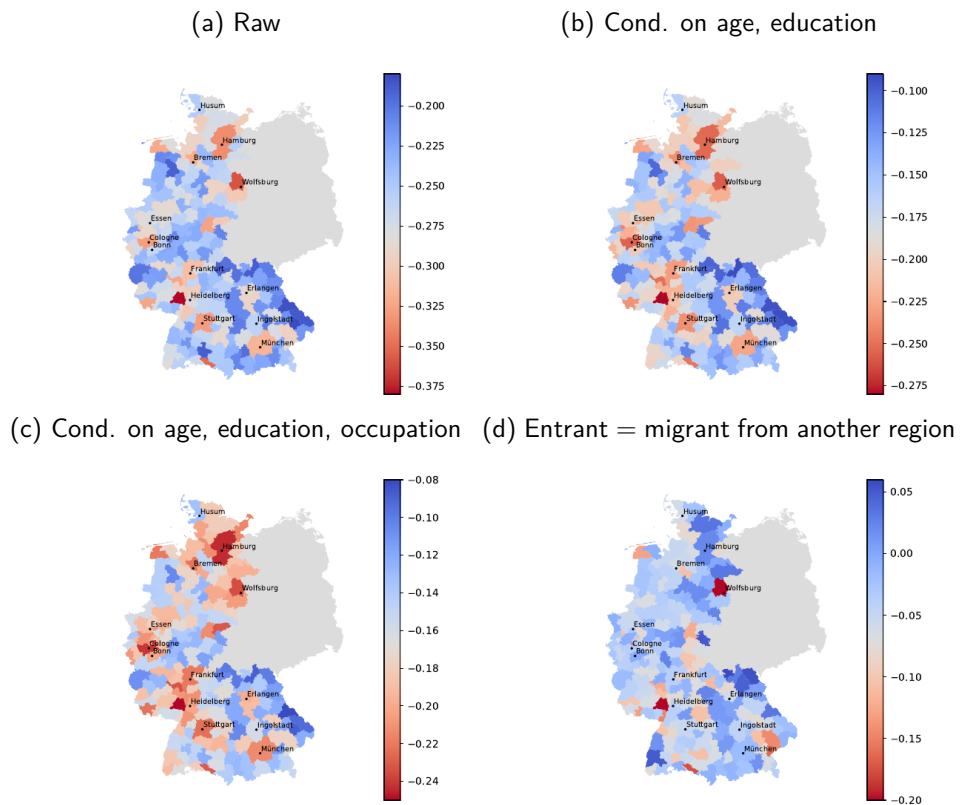
B.4 Entrants' wages

Figure 3.17: Wages of regional entrants minus wages of incumbents



Notes: The vertical axis in Panel 3.17a shows the average wage of an entrant to a region relative to the average wage of incumbents in that region. The vertical axis in Panel 3.17b shows the (relative) residual wage of entrants after a regression on age and education dummies. Panel 3.17c repeats this and additionally adds occupation dummies to the regression. The vertical axis in Panel 3.17d shows the average wage of a mover from another local labor market relative to the average wage of incumbents in the destination region. Averages are taken across years 1985 until 2010. The horizontal axes depict the change in log employment between 1985 and 2010. One bubble represents one of 182 West German local labor markets aggregated from political districts as suggested by Kropp and Schwengler (2011). Bubble size corresponds to the number of employed workers in a market averaged across years 1985 to 2010. The four groups show a classification of local labor markets into northern and southern, rural and urban areas (see Appendix Figure 3.10). Southern areas are located in Baden-Wuerttemberg and Bavaria. The classification of local labor markets into rural and urban is based on BBSR (2014). Regression lines across all areas (black) and within the four broad groups (colored) are weighted by the number of employed workers.

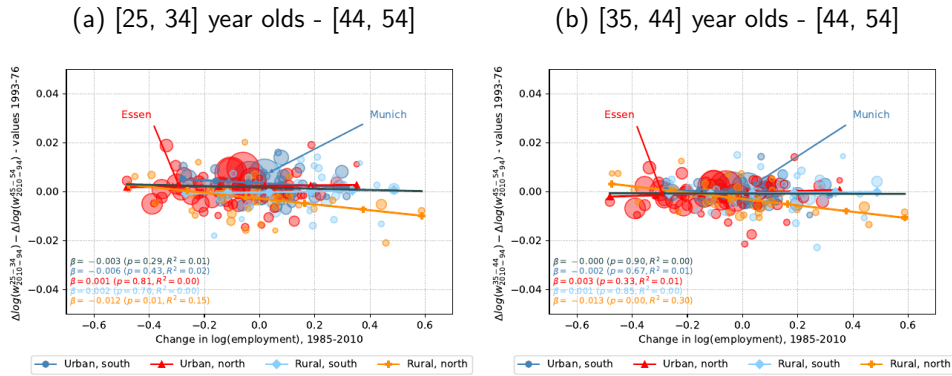
Figure 3.18: Entrants' relative wages across space



Notes: The map in Panel 3.18a shows the average wage of an entrant to a region relative to the average wage of incumbents in that region. The map in Panel 3.18a shows the (relative) residual wage of entrants after a regression on age and education dummies. Panel 3.18a repeats this and additionally adds occupation dummies to the regression. The map in Panel 3.18a shows the average wage of a mover from another local labor market relative to the average wage of incumbents in the destination region. Averages are taken across years 1985 until 2010. Assignment of political districts to West German local labor markets is based on commuting flows and follows Kropp and Schwengler (2011). East Germany is left out of the analysis.

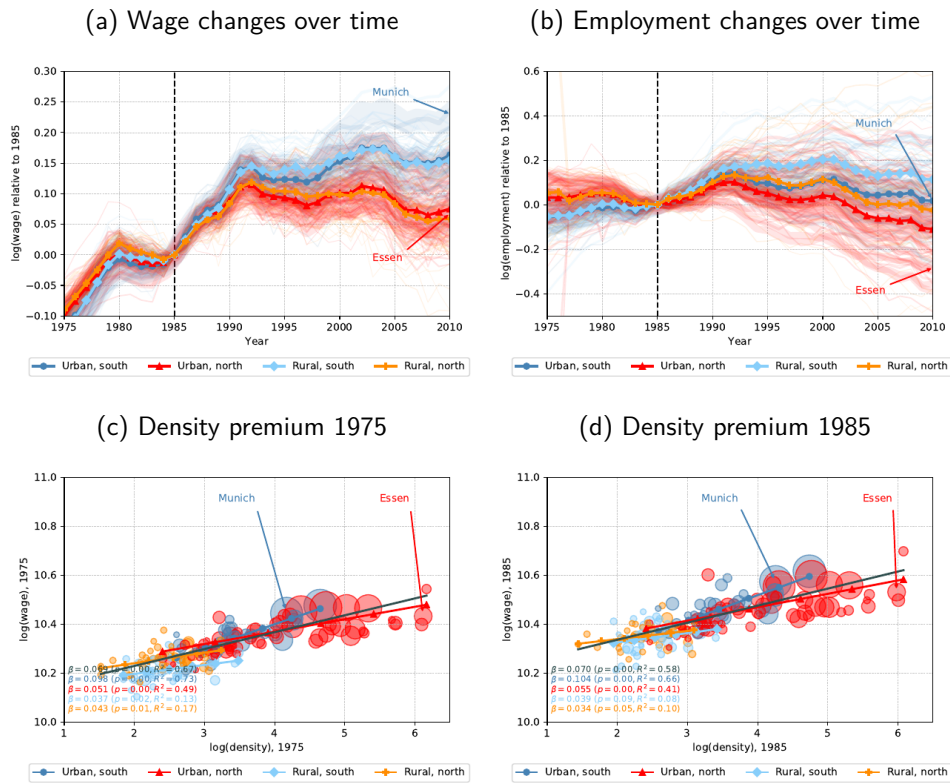
C Additional Results for Section 3.4.2

Figure 3.19: Wage growth between age groups



Notes: The figures show a triple difference-in-difference result: how much has wage growth of young (Figure 3.19a) and medium old (Figure 3.19b) workers relative to the wage growth of old workers changed after 1993 relative to 1993 and before? The horizontal axes depict the change in log employment between 1985 and 2010. One bubble represents one of 182 West German local labor markets aggregated from political districts as suggested by Kropp and Schwengler (2011). Bubble size corresponds to the number of employed workers in a market averaged across years 1985 to 2010. The four groups show a classification of local labor markets into northern and southern, rural and urban areas (see Appendix Figure 3.10). Southern areas are located in Baden-Wuerttemberg and Bavaria. The classification of local labor markets into rural and urban is based on BBSR (2014). Regression lines across all areas (black) and within the four broad groups (colored) are weighted by the number of employed workers.

Figure 3.20: Changes in employment and wages including base period




Notes: Figure 3.20a shows the change in log wages within local labor markets over time. Figure 3.20b shows the change in log employment. Shaded lines in the background represent 182 West German detailed local labor markets. The four groups show a classification of local labor markets into northern and southern, rural and urban areas (see Appendix Figure 3.10). Southern areas are located in Baden-Wuerttemberg and Bavaria. The classification of local labor markets into rural and urban is based on BBSR (2014). The thickness of a shaded background line corresponds to the number of employed workers in a local labor market averaged across years 1985 until 2010. Figure 3.20c shows average wages in 1975 plotted against market density defined as number of employed workers divided by area in square kilometers. Figure 3.20d repeats this for 1985. One bubble represents one of 182 West German local labor markets aggregated from political districts as suggested by Kropp and Schwengler (2011). Bubble size corresponds to the number of employed workers in a market averaged across years 1985 to 2010. The four groups show a classification of local labor markets into northern and southern, rural and urban areas (see Appendix Figure 3.10). Southern areas are located in Baden-Wuerttemberg and Bavaria. The classification of local labor markets into rural and urban is based on BBSR (2014). Regression lines across all areas (black) and within the four broad groups (colored) are weighted by the number of employed workers.

D Additional Results for Section 3.5

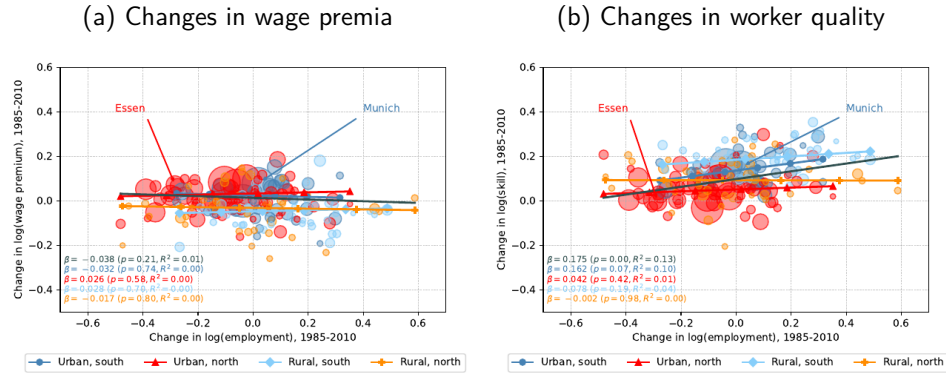
Figure 3.21: Skill accumulation

(a) Dense places have higher skill accumulation (b) Skill accumulation across space



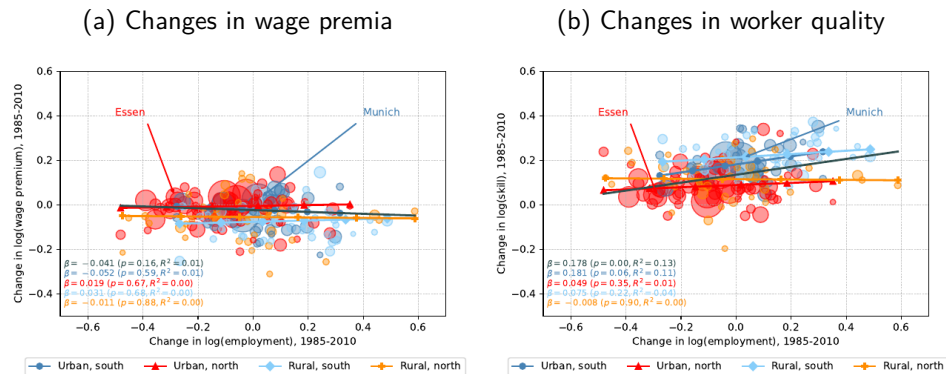
Notes: Figure 3.21a shows estimated skill accumulation rates $\Gamma_{l,l}$ for a hypothetical market stayer at age 54. OLS estimates as described by Equation (3.9). The horizontal axis depicts log employment density 1985 defined as number of employed workers in a market divided by area in square kilometers. One bubble represents one of 182 West German local labor markets aggregated from political districts as suggested by Kropp and Schwengler (2011). Bubble size corresponds to the number of employed workers in a market averaged across years 1985 to 2010. The four groups show a classification of local labor markets into northern and southern, rural and urban areas (see Appendix Figure 3.10). Southern areas are located in Baden-Wuerttemberg and Bavaria. The classification of local labor markets into rural and urban is based on BBSR (2014). Regression lines across all areas (black) and within the four broad groups (colored) are weighted by the number of employed workers. Figure 3.21b plots estimated skill accumulation rates across Germany. Assignment of political districts to West German local labor markets is based on commuting flows and follows Kropp and Schwengler (2011). East Germany is left out of the analysis.

Figure 3.22: Omitting controls for changes in amenities



Notes: Figure 3.22a shows estimated changes in wage premia between 1985 and 2010. OLS estimates as described by Equation (3.9) but omitting controls for changing amenities. Figure 3.22b shows estimated changes in worker quality as described by Equation (3.11) again omitting controls for changing amenities. The horizontal axes depict the change in log employment between 1985 and 2010. One bubble represents one of 182 West German local labor markets aggregated from political districts as suggested by Kropp and Schwengler (2011). Bubble size corresponds to the number of employed workers in a market averaged across years 1985 to 2010. The four groups show a classification of local labor markets into northern and southern, rural and urban areas (see Appendix Figure 3.10). Southern areas are located in Baden-Wuerttemberg and Bavaria. The classification of local labor markets into rural and urban is based on BBSR (2014). Regression lines across all areas (black) and within the four broad groups (colored) are weighted by the number of employed workers.

Figure 3.23: Education dependent skill accumulation function




Notes: Figure 3.23a shows estimated changes in wage premia between 1985 and 2010. OLS estimates as described by Equation (3.9) but omitting controls for changing amenities. Additionally, speed of skill accumulation depends on three education groups. University degree: technical college or university; Medium educated: Abitur or apprenticeship training; Low educated: without postsecondary education or missing. Figure 3.23b shows estimated changes in worker quality as described by Equation (3.11) again omitting controls for changing amenities. The horizontal axes depict the change in log employment between 1985 and 2010. One bubble represents one of 182 West German local labor markets aggregated from political districts as suggested by Kropp and Schwengler (2011). Bubble size corresponds to the number of employed workers in a market averaged across years 1985 to 2010. The four groups show a classification of local labor markets into northern and southern, rural and urban areas (see Appendix Figure 3.10). Southern areas are located in Baden-Wuerttemberg and Bavaria. The classification of local labor markets into rural and urban is based on BBSR (2014). Regression lines across all areas (black) and within the four broad groups (colored) are weighted by the number of employed workers.

Figure 3.24: Occupation dependent skill accumulation function

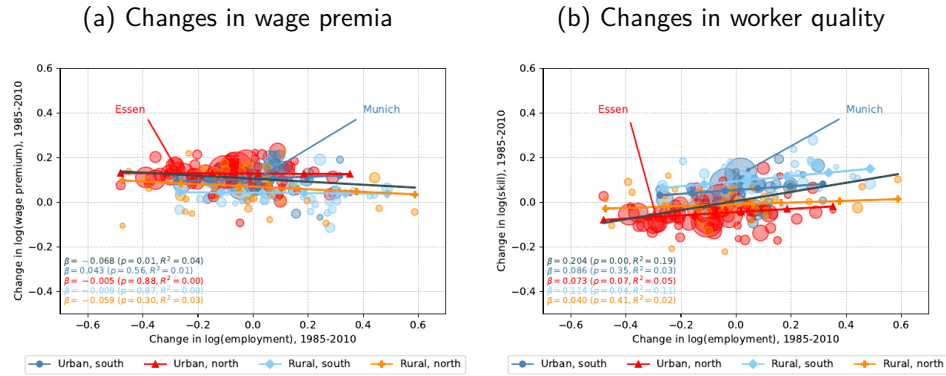
(a) Changes in wage premia

(b) Changes in worker quality



Notes: Figure 3.24a shows estimated changes in wage premia between 1985 and 2010. OLS estimates as described by Equation (3.9) but omitting controls for changing amenities. Additionally, speed of skill accumulation depends on four occupation groups. Mgr-Prof-Tech: managers, professionals, technicians; Prod-Op-Crafts: production, operators, craftsmen; Srv-Care: services and care. Classification of detailed occupations follows Acemoglu and Autor (2011), see Böhm et al. (2019) for the exact mapping. Figure 3.24b shows estimated changes in worker quality as described by Equation (3.11) again omitting controls for changing amenities. The horizontal axes depict the change in log employment between 1985 and 2010. One bubble represents one of 182 West German local labor markets aggregated from political districts as suggested by Kropp and Schwengler (2011). Bubble size corresponds to the number of employed workers in a market averaged across years 1985 to 2010. The four groups show a classification of local labor markets into northern and southern, rural and urban areas (see Appendix Figure 3.10). Southern areas are located in Baden-Wuerttemberg and Bavaria. The classification of local labor markets into rural and urban is based on BBSR (2014). Regression lines across all areas (black) and within the four broad groups (colored) are weighted by the number of employed workers.


Figure 3.25: Fixed effects estimation



Notes: Figure 3.25a shows estimated changes in wage premia between 1985 and 2010. Results were obtained from a fixed effects estimation including a separate worker-region fixed effect each time a worker revisits a local labor market similar to Cortes (2016). Wage premia are identified from year-region fixed effects. Region times detailed age dummies control for life cycle skill accumulation. Figure 3.25b shows estimated changes in worker quality as described by Equation (3.11) using wage premia estimated as in Figure 3.25a. The horizontal axes depict the change in log employment between 1985 and 2010. One bubble represents one of 182 West German local labor markets aggregated from political districts as suggested by Kropp and Schwengler (2011). Bubble size corresponds to the number of employed workers in a market averaged across years 1985 to 2010. The four groups show a classification of local labor markets into northern and southern, rural and urban areas (see Appendix Figure 3.10). Southern areas are located in Baden-Wuerttemberg and Bavaria. The classification of local labor markets into rural and urban is based on BBSR (2014). Regression lines across all areas (black) and within the four broad groups (colored) are weighted by the number of employed workers.

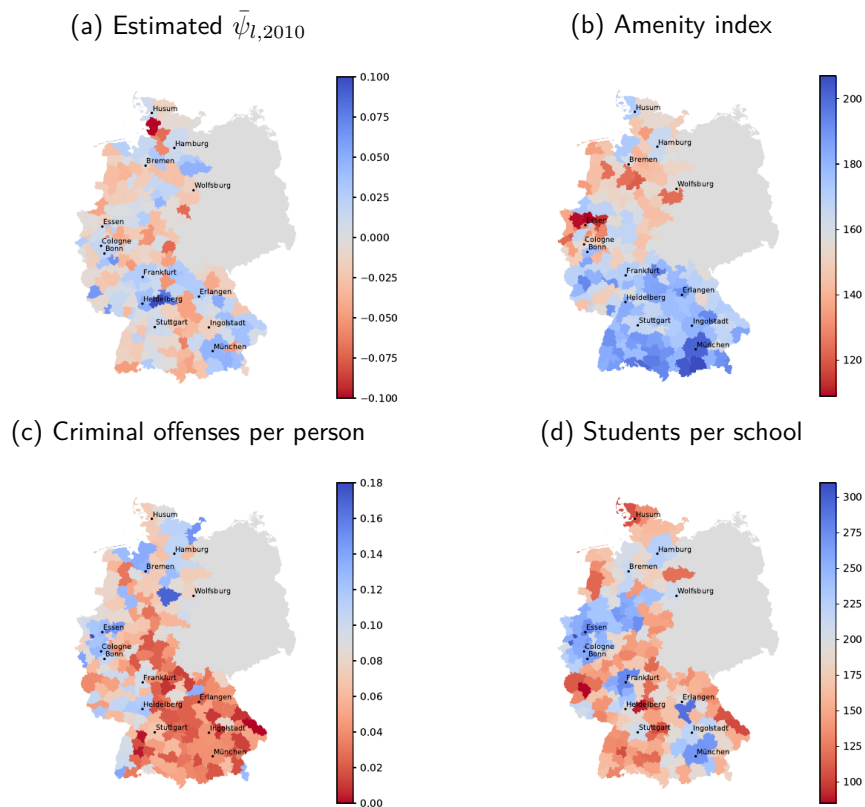
Figure 3.26: Reweighted skill changes, education and industry

(a) Match 1985's education structure (b) Match 1985's industry structure

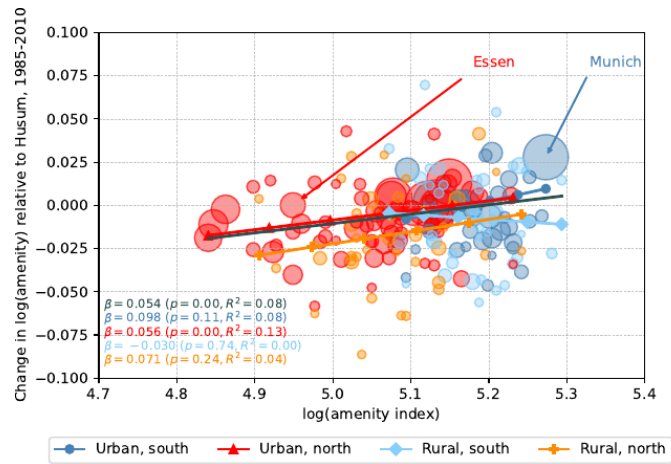


Notes: Figure 3.26a shows estimated changes in worker quality as described by Equation (3.11). Observations were reweighted following DiNardo et al. (1996) to eliminate differential changes in the education structure across local labor markets. Figure 3.26b repeats this by reweighting observations to eliminate differential changes in the industry structure across local labor markets. The horizontal axes depict the change in log employment between 1985 and 2010. One bubble represents one of 182 West German local labor markets aggregated from political districts as suggested by Kropp and Schwengler (2011). Bubble size corresponds to the number of employed workers in a market averaged across years 1985 to 2010. The four groups show a classification of local labor markets into northern and southern, rural and urban areas (see Appendix Figure 3.10). Southern areas are located in Baden-Wuerttemberg and Bavaria. The classification of local labor markets into rural and urban is based on BBSR (2014). Regression lines across all areas (black) and within the four broad groups (colored) are weighted by the number of employed workers.

Figure 3.27: Amenities across space



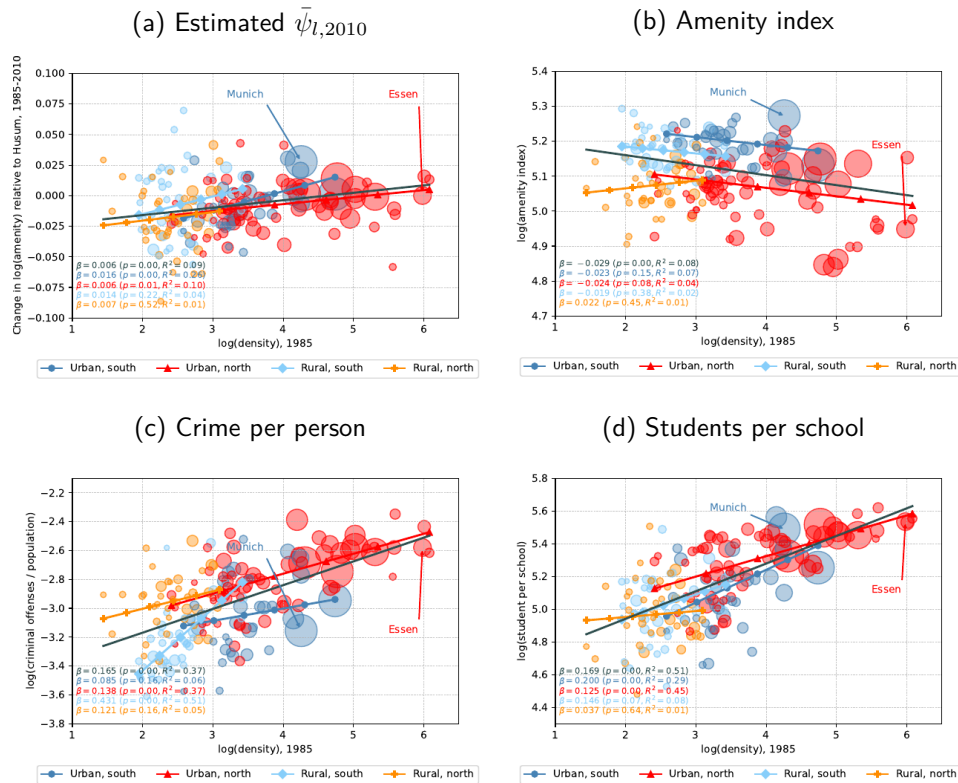
Notes: The maps show the distribution of estimated amenities $\bar{\psi}_{i,2010}$ according to Equation (3.9) as well as collected amenity proxies across space. The amenity index was computed by prognos (2018) combining information on 53 amenity values ranging from sunshine duration to restaurant and physician density. Data on criminal offenses comes from Bundeskriminalamt (2018). The measure was computed as offenses divided by population collected from Statistische Ämter des Bundes und der Länder (2017a). Information on students per school was retrieved from Statistische Ämter des Bundes und der Länder (2017b). Assignment of political districts to West German local labor markets is based on commuting flows and follows Kropp and Schwengler (2011). East Germany is left out of the analysis.

Figure 3.28: Estimated amenities $\bar{\psi}_{l,2010}$ and the external amenity index

Notes: The vertical axis shows estimated amenity values $\bar{\psi}_{l,2010}$ according to Equation (3.9). The horizontal axis shows an amenity index computed by prognos (2018) combining information on 53 amenity values ranging from sunshine duration to restaurant and physician density. One bubble represents one of 182 West German local labor markets aggregated from political districts as suggested by Kropp and Schwengler (2011). Bubble size corresponds to the number of employed workers in a market averaged across years 1985 to 2010. The four groups show a classification of local labor markets into northern and southern, rural and urban areas (see Appendix Figure 3.10). Southern areas are located in Baden-Wuerttemberg and Bavaria. The classification of local labor markets into rural and urban is based on BBSR (2014). Regression lines across all areas (black) and within the four broad groups (colored) are weighted by the number of employed workers.

E Additional Results for Section 3.6


Figure 3.29: Relation between amenities and density



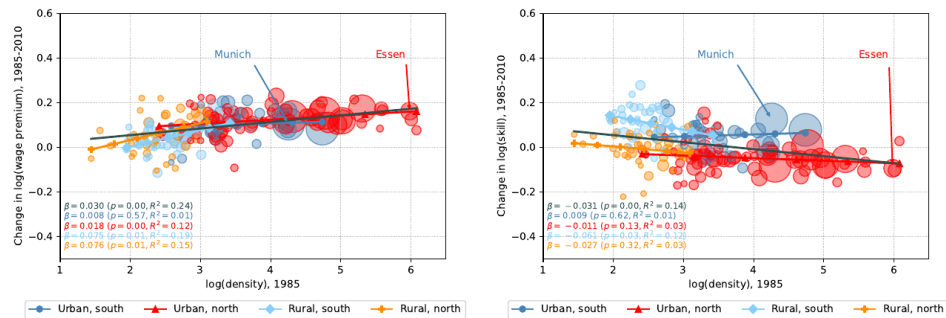
Notes: The vertical axes show estimated amenities $\bar{\psi}_{l,2010}$ according to Equation (3.9) (Figure 3.29a) as well as collected amenity proxies. The horizontal axes depicts log employment density 1985 defined as number of employed workers in a market divided by area in square kilometers. The amenity index was computed by prognos (2018) combining information on 53 amenity values ranging from sunshine duration to restaurant and physician density. Data on criminal offenses comes from Bundeskriminalamt (2018). The measure was computed as offenses divided by population collected from Statistische Ämter des Bundes und der Länder (2017a). Information on students per school was retrieved from Statistische Ämter des Bundes und der Länder (2017b). One bubble represents one of 182 West German local labor markets aggregated from political districts as suggested by Kropp and Schwengler (2011). Bubble size corresponds to the number of employed workers in a market averaged across years 1985 to 2010. The four groups show a classification of local labor markets into northern and southern, rural and urban areas (see Appendix Figure 3.10). Southern areas are located in Baden-Wuerttemberg and Bavaria. The classification of local labor markets into rural and urban is based on BBSR (2014). Regression lines across all areas (black) and within the four broad groups (colored) are weighted by the number of employed workers.

Figure 3.30: Wage premia vs density, omitting amenities, fixed effects

(a) Changes in wage premia, no amenities (b) Changes in worker quality, no amenities



(c) Changes in wage premia, fixed effects (d) Changes in worker quality, fixed effects



Notes: Figure 3.30a shows estimated changes in wage premia between 1985 and 2010. OLS estimates as described by Equation (3.9) but omitting controls for changing amenities. Figure 3.30b shows estimated changes in worker quality as described by Equation (3.11) again omitting controls for changing amenities. Figure 3.30c shows estimated changes in wage premia by means of a fixed effects estimation including a separate worker-region fixed effect each time a worker revisits a local labor market similar to Cortes (2016). Wage premia are identified from year-region fixed effects. Region times detailed age dummies control for life cycle skill accumulation. Figure 3.30d shows estimated changes in worker quality as described by Equation (3.11) using wage premia estimated as in Figure 3.30c. The horizontal axes depicts log employment density 1985 defined as number of employed workers in a market divided by area in square kilometers. One bubble represents one of 182 West German local labor markets aggregated from political districts as suggested by Kropp and Schwengler (2011). Bubble size corresponds to the number of employed workers in a market averaged across years 1985 to 2010. The four groups show a classification of local labor markets into northern and southern, rural and urban areas (see Appendix Figure 3.10). Southern areas are located in Baden-Wuerttemberg and Bavaria. The classification of local labor markets into rural and urban is based on BBSR (2014). Regression lines across all areas (black) and within the four broad groups (colored) are weighted by the number of employed workers.

References

- Abowd, John M, Francis Kramarz, and David N Margolis (1999). "High Wage Workers and High Wage Firms". In: *Econometrica* 67.2, pages 251–333.
- Acemoglu, Daron and David H. Autor (2011). "Skills, Tasks and Technologies: Implications for Employment and Earnings". In: *Handbook of Labor Economics*, vol. 4 Part B. Edited by David Card and Orley Ashenfelter. Elsevier. Chapter 12, pages 1043–1171.
- Amior, Michael and Alan Manning (2018). "The Persistence of Local Joblessness". In: *American Economic Review* 108.7, pages 1942–1970.
- Arntz, Melanie (2010). "What Attracts Human Capital? Understanding the Skill Composition of Interregional Job Matches in Germany". In: *Regional Studies* 44.4, pages 423–441.
- Autor, David (2019). "Work of the Past, Work of the Future". In: *Richard T. Ely Lecture given at Annual Meeting of American Economics Association, Atlanta, GA, January*. Volume 4.
- Autor, David, David Dorn, Gordon Hanson, and Kaveh Majlesi (2016). "Importing Political Polarization? The Electoral Consequences of Rising Trade Exposure". Working Paper.
- Bárány, Zsófia and Christian Siegel (2018). "Job Polarization and Structural Change". In: *American Economic Journal: Macroeconomics* 10.1, pages 59–89.
- Baumert, Jürgen, Cordula Artelt, Eckhard Klieme, Michael Neubrand, Manfred Prenzel, Ulrich Schiefele, Wolfgang Schneider, Klaus-Jürgen Tillmann, and Manfred Weiß (2013). *PISA 2000—Die Länder der Bundesrepublik Deutschland im Vergleich*. Springer-Verlag.
- BBSR (2014). *Bundesinstitut für Bau-, Stadt- und Raumforschung: Raumabgrenzungen*. URL: https://www.bbsr.bund.de/BBSR/DE/Raumbearbeitung/Downloads/downloads_node.html.
- Berry, Christopher R. and Edward L. Glaeser (2005). "The divergence of human capital levels across cities". In: *Papers in Regional Science* 84.3, pages 407–444.
- Blanchard, Olivier Jean and Lawrence F. Katz (1992). "Regional Evolutions". In: *Brookings Papers on Economic Activity* 1992.1, pages 1–75.
- Bleakley, Hoyt and Jeffrey Lin (2012). "Thick-Market Effects and Churning in the Labor Market: Evidence from US Cities". In: *Journal of Urban Economics* 72.2, pages 87–103.
- Böhm, Michael, Hans-Martin von Gaudecker, and Felix Schran (2019). "Occupation Growth, Skill Prices, and Wage Inequality". Working Paper.
- Bound, John and Harry J. Holzer (2000). "Demand Shifts, Population Adjustments, and Labor Market Outcomes During the 1980s". In: *Journal of Labor Economics* 18.1, pages 20–54.
- Bundeskriminalamt (2018). *Polizeiliche Kriminalstatistik, PKS 2018 - BKA (Länder-, Kreis- und Städte) Übersicht Falltabellen*. URL: <https://www.bka.de>.

- de/DE/AktuelleInformationen/StatistikenLagebilder/PolizeilicheKriminalstatistik/PKS2018/BKATabellen/bkaTabellenLaenderKreiseStaedteFaelle.html?nn=108686.
- Card, David, Jörg Heining, and Patrick Kline (2013). "Workplace Heterogeneity and the Rise of West German Wage Inequality". In: *Quarterly Journal of Economics* 128.3, pages 967–1015.
- Combes, Pierre-Philippe, Gilles Duranton, and Laurent Gobillon (2008). "Spatial Wage Disparities: Sorting Matters!" In: *Journal of Urban Economics* 63.2, pages 723–742.
- Combes, Pierre-Philippe, Gilles Duranton, Laurent Gobillon, and Sébastien Roux (2012). "Sorting and Local Wage and Skill Distributions in France". In: *Regional Science and Urban Economics* 42.6, pages 913–930.
- Cortes, Guido Matias (2016). "Where Have the Middle-Wage Workers Gone? A Study of Polarization Using Panel Data". In: *Journal of Labor Economics* 34.1, pages 63–105.
- Czernich, Nina, Oliver Falck, Tobias Kretschmer, and Ludger Woessmann (2011). "Broadband Infrastructure and Economic Growth". In: *The Economic Journal* 121.552, pages 505–532.
- D'Costa, Sabine and Henry G Overman (2014). "The Urban Wage Growth Premium: Sorting or Learning?" In: *Regional Science and Urban Economics* 48, pages 168–179.
- Dahl, Gordon B. (2002). "Mobility and the Return to Education: Testing a Roy Model with Multiple Markets". In: *Econometrica* 70.6, pages 2367–2420.
- Dauth, Wolfgang, Sebastian Findeisen, Enrico Moretti, and Jens Südekum (2019). "Matching in Cities". IZA Discussion Paper No. 12278.
- Dauth, Wolfgang, Sebastian Findeisen, and Jens Südekum (2017). "Trade and Manufacturing Jobs in Germany". In: *American Economic Review: Papers & Proceedings* 107.5, pages 337–342.
- De la Roca, Jorge and Diego Puga (2017). "Learning by Working in Big Cities". In: *Review of Economic Studies* 84.1, pages 106–142.
- Diamond, Rebecca (March 2016). "The Determinants and Welfare Implications of US Workers' Diverging Location Choices by Skill: 1980-2000". In: *American Economic Review* 106.3, pages 479–524.
- DiNardo, John, Nicole M. Fortin, and Thomas Lemieux (1996). "Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semiparametric Approach". In: *Econometrica* 64.5, pages 1001–1044.
- Duranton, Gilles (2007). "Urban Evolutions: The Fast, the Slow, and the Still". In: *American Economic Review* 97.1.
- Duranton, Gilles and Diego Puga (2004). "Micro-foundations of Urban Agglomeration Economies". In: *Handbook of Regional and Urban Economics*. Volume 4. Elsevier, pages 2063–2117.
- Dustmann, Christian and Albrecht Glitz (2015). "How do Industries and Firms Respond to Changes in Local Labor Supply?" In: *Journal of Labor Economics* 33.3, pages 711–750.

- Dustmann, Christian, Johannes Ludsteck, and Uta Schönberg (2009). "Revisiting the German wage structure". In: *Quarterly Journal of Economics* 124.2, pages 843–881.
- Eeckhout, Jan, Roberto Pinheiro, and Kurt Schmidheiny (2014). "Spatial Sorting". In: *Journal of Political Economy* 122.3, pages 554–620.
- Findeisen, Sebastian and Jens Südekum (2008). "Industry Churning and the Evolution of Cities: Evidence for Germany". In: *Journal of Urban Economics* 64.2, pages 326–339.
- Fitzenberger, Bernd, Anderonke Osikominu, and Robert Völter (2006). "Imputation Rules to Improve the Education Variable in the IAB Employment Sub-sample". In: *Schmollers Jahrbuch (Journal of the Applied Social Sciences)* 126, pages 405–436.
- Ganzer, Andreas, Alexandra Schmucker, Philipp vom Berge, and Anja Wurdack (2017). *Sample of Integrated Labour Market Biographies – Regional File 1975-2014 (SIAB-R 7514)*. Technical report 1. FDZ–Datenreport.
- Glaeser, Edward L (1999). "Learning in Cities". In: *Journal of Urban Economics* 46.2, pages 254–277.
- Glaeser, Edward L (2005). "Reinventing Boston: 1630–2003". In: *Journal of Economic Geography* 5.2, pages 119–153.
- Glaeser, Edward L. (2008). *Cities, Agglomeration, and Spatial Equilibrium*. Cambridge, UK: Oxford University Press.
- Glaeser, Edward L. (2011). *Triumph of the City: How Our Greatest Invention Makes Us Richer, Smarter, Greener, Healthier and Happier*. New York, NY: Penguin Press.
- Glaeser, Edward L. and Joshua D. Gottlieb (2008). "The Economics of Place-Making Policies". In: *Brookings Papers on Economic Activity* 1, pages 155–253.
- Glaeser, Edward L., Jed Kolko, and Albert Saiz (2001). "Consumer City". In: *Journal of Economic Geography* 1.1, pages 27–50.
- Glaeser, Edward L. and David C. Maré (2001). "Cities and Skills". In: *Journal of Labor Economics* 19.2, pages 316–342.
- Heckman, James J., Lance Lochner, and Christopher Taber (1998). "Explaining Rising Wage Inequality: Explorations with a Dynamic General Equilibrium Model of Labor Earnings with Heterogeneous Agents". In: *Review of Economic Dynamics* 1, pages 1–58.
- Hsieh, Chang-Tai, Erik Hurst, Charles I. Jones, and Peter J. Klenow (2019). "The Allocation of Talent and U.S. Economic Growth". In: *Econometrica* 87.5, pages 1439–1474.
- Kawka, Rupert, Sabine Beisswenger, Gabriele Costa, Heike Kemmerling, Sarah Müller, Thomas Pütz, Helena Schmidt, Stefan Schmidt, and Marisa Trimborn (2009). *Regionaler Preisindex, Berichte Band 30*. Technical report. URL: https://www.bbsr.bund.de/BBSR/DE/Veroeffentlichungen/Abgeschlossen/Berichte/2009_2010/Bd30RegPreisindex.html. Bonn: Bundesinstitut für Bau-, Stadt- und Raumforschung.

- Kline, Patrick and Enrico Moretti (2014). "People, Places, and Public Policy: Some Simple Welfare Economics of Local Economic Development Programs". In: *Annual Review of Economics* 6.1, pages 629–662.
- Kropp, Per and Barbara Schwengler (2011). "Abgrenzung von Arbeitsmarktregionen - ein Methodenvorschlag". In: *Raumforschung und Raumordnung* 1, pages 45–62.
- Manning, Alan and Barbara Petrongolo (2017). "How Local are Labor Markets? Evidence from a Spatial Job Search Model". In: *American Economic Review* 107.10, pages 2877–2907.
- Moretti, Enrico (2012). *The New Geography of Jobs*. New York: Houghton Mifflin Harcourt.
- Moretti, Enrico (2013). "Real Wage Inequality". In: *American Economic Journal: Applied Economics* 5.1, pages 65–103.
- prognos (2018). *Hintergrund zur Großen ZDF-Deutschland-Studie*. Technical report. Berlin: prognos AG. URL: http://deutschlandstudie.prognos.com/fileadmin/Dokumente/kdjvbkjsdbvlasdjvlihsdfdaihfaihdgh1h1isdfhvai4/Hintergrund-PDF_Deutschlandstudie_Mai_2018.
- Roback, Jennifer (1988). "Wages, Rents, and Amenities: Differences Among Workers and Regions". In: *Economic Inquiry* 26.1, pages 23–41.
- Rosen, Sherwin (1986). "Chapter 12: The Theory of Equalizing Differences". In: edited by Orley C. Ashenfelter and Richard Layard. Volume 1. *Handbook of Labor Economics Supplement C*. Elsevier, pages 641–692.
- Roy, A. D. (1951). "Some Thoughts on the Distribution of Earnings". In: *Oxford Economic Papers* 3.2, pages 135–146.
- Sorkin, Isaac (2018). "Ranking Firms Using Revealed Preference". In: *Quarterly Journal of Economics* 133.3, pages 1331–1393.
- Statistische Ämter des Bundes und der Länder (2017a). *Regionaldatenbank Deutschland: Fortschreibung des Bevölkerungsstandes*. URL: <https://www.regionalstatistik.de/genesis/online/data;jsessionid=63B2B49818074223523B1E3FABD6039A.reg2?operation=previous&levelindex=3&levelid=1558971972857&levelid=1558971948778&step=2>.
- Statistische Ämter des Bundes und der Länder (2017b). *Regionaldatenbank Deutschland: Statistik der allgemeinbildenden Schulen*. URL: <https://www.regionalstatistik.de/genesis/online/data;jsessionid=69E491BC7F859D12B21C590C0D19D109.reg2?operation=abrufetabelleAbrufen&selectionname=21111-01-03-4&levelindex=1&levelid=1558971361255&index=1>.
- Südekum, Jens (2008). "Convergence of the Skill Composition Across German Regions". In: *Regional Science and Urban Economics* 38.2, pages 148–159.
- Wheaton, William C. (1979). "Metropolitan Growth, Unemployment, and Inter-regional Factor Mobility". In: *Interregional Movements and Regional Growth*. Edited by William C. Wheaton. Urban Institute, pages 237–253.

- Wheeler, Christopher H (2006). "Cities and the Growth of Wages Among Young Workers: Evidence from the NLSY". In: *Journal of Urban Economics* 60.2, pages 162–184.
- Yamaguchi, Shintaro (2018). "Changes in Returns to Task-Specific Skills and Gender Wage Gap". In: *Journal of Human Resources* 53.1, pages 32–70.

Lebenslauf

Geboren am 22. Juni 1990 in Winterberg, wuchs ich in Oberkirchen-Inderlenne auf. Im Jahr 2009 erlangte ich am Gymnasium Schmallebenberg die allgemeine Hochschulreife. Im Sommersemester 2010 nahm ich das Studium der Volkswirtschaftslehre an der Universität Bonn auf, das ich im Dezember 2013 mit dem Abschluss Bachelor of Science beendete. Anschließend entschied ich mich für ein Promotionsstudium an der Bonn Graduate School of Economics. Im Laufe des Promotionsstudiums verbrachte ich im Herbst 2017 ein Semester am University College London. Den Masterabschluss erlangte ich im Dezember 2017.